

Branching random tessellations with interaction: a thermodynamic view

Hans-Otto Georgii*, Tomasz Schreiber† and Christoph Thäle‡§

Abstract

A branching random tessellation (BRT) is a stochastic process that transforms a coarse initial tessellation of \mathbb{R}^d into a finer tessellation by means of random cell divisions in continuous time. This concept generalises the so-called STIT tessellations, for which all cells split up independently of each other. Here we allow the cells to interact, in that the division rule for each cell may depend on the structure of the surrounding tessellation. Moreover, we consider coloured tessellations, for which each cell is marked with an internal property, called its colour. Under a suitable condition, the cell interaction of a BRT can be specified by a measure kernel that determines the division rules of all cells and gives rise to a Gibbsian characterisation of BRTs. For translation invariant BRTs, we introduce an ‘inner’ entropy density relative to a STIT tessellation. Together with an inner energy density for a given ‘moderate’ division kernel, this leads to a variational principle for BRTs with this prescribed kernel, and further to an existence result for such BRTs.

Keywords. Stochastic geometry, branching tessellation, coloured tessellation, STIT tessellation; Gibbs measure, relative entropy, free energy, variational principle.

MSC (2010). Primary: 60D05, 60K35; Secondary: 28D20, 60G55, 82B21.

This article is dedicated to the memory of the second author, Tomasz Schreiber. Shortly before his untimely death of cancer at the age of 35, he entrusted to the first author a draft manuscript that introduced the general setup of this paper and initiated a study of relative entropy densities for BRTs. This paper tries to realise his intentions with various modifications of concepts and results.

1 Introduction

A central object of stochastic geometry and spatial stochastics are tessellations of \mathbb{R}^d (with $d \geq 1$), i.e. locally finite families of d -dimensional convex polytopes that cover \mathbb{R}^d and have pairwise disjoint interiors. They are used in many practical applications. For example, random tessellations serve as models for cellular or polycrystalline materials, plant cells or influence zones, for instance in the modelling of telecommunication networks or animal territories; see [14, 24] for an overview.

The standard random tessellations usually considered in the literature are the Poisson hyperplane tessellations, the Poisson-Voronoi and the Poisson-Delaunay tessellations; cf. [17] for definitions. These have the property of being facet-to-facet (or side-to-side in the planar case), which is to say that the intersection of any two of its cells is either empty or a common face of both cells. However, there are numerous applications for which models of this kind are inappropriate, for example network models

*Mathematisches Institut der Universität München, Theresienstraße 39, 80333 München, Germany. E-mail: georgii@math.lmu.de

†Faculty of Mathematics and Computer Science, Nicolaus Copernicus University, Toruń, Poland.

‡Fakultät für Mathematik der Ruhr-Universität Bochum, Universitätsstraße 150, 44801 Bochum, Germany. E-mail: christoph.thaele@rub.de

§Partially supported by the German Research Foundation (DFG) via SFB/TR-12 “Symmetries and Universality in Mesoscopic Systems”.

for telecommunication systems or models for crack structures in geology. Hence, there is a growing demand for mathematically tractable models of non-facet-to-facet tessellations, which may serve as idealised reference models. Only some years ago, the class of iteration-stable random tessellations (called STIT tessellations for short) was introduced by Nagel and Weiß in [13]. These tessellations are constructed by means of a random process of cell division and have attracted considerable interest because of its analytical tractability; see for example [15], [18]–[23] or [25].

Our objects of study here are generalisations of the STIT models. On the one hand, we consider coloured tessellations, for which each cell is equipped with an individual colour. For example, the colour of a cell could represent its nutrient content, its genotype, age, or whatever else might be relevant to describe the state of a cell. (In a different context, coloured tessellations have been studied by Arak and Surgailis [1, 2], for example.) On the other hand, and more importantly, we allow for an interaction of cells during their division process. Informally, the procedure is the following. At time zero, one starts with an initial random tessellation of \mathbb{R}^d into coloured cells. Each cell lives for a random time, which is determined by an interactive competition of cells. Namely, the survival rate of a cell c at any time $s > 0$ may not only depend on the cell's geometry and colour, but in fact on the whole tessellation including its past evolution. When the lifetime has run out, a hyperplane with coloured half-spaces is chosen randomly according to some rule that may again depend on the cell's geometry, colour and the past evolution of the surrounding tessellation, and is used to cut c into two polyhedral sub-cells c^+ and c^- , which inherit their colours from the respective half-spaces of the cutting hyperplane. The daughter cells c^+ and c^- then replace c in the collective division game, which is continued until time 1, say. The resulting tessellation of \mathbb{R}^d at a deterministic time $s > 0$ is denoted by T_s , and the tessellation-valued stochastic process $(T_s)_{s \in [0,1]}$ is what we call a *branching random tessellation* or BRT for short.

In the special case when (i) the distribution of lifetimes is exponential with parameter proportional to the mean width of the cells, and (ii) the bicoloured hyperplanes are chosen at random according to the motion-invariant hyperplane measure and some reference measure on the colour space, $(T_s)_{s \in [0,1]}$ is a coloured STIT tessellation of \mathbb{R}^d and its distribution is invariant under rigid motions whenever so is the initial random tessellation. The coloured STIT tessellations play an important rôle in the background of our theory, but our focus is on cell-splitting mechanisms with interaction as described above. Different tessellation models with cell interaction, namely Delaunay or Voronoi tessellations of Gibbsian type, are discussed in [3, 5, 6].

The main results of this paper are the following.

- First, we investigate the behaviour of general BRTs, typically denoted by \mathbf{P} . Under a condition of local absolute continuity (LAC), we establish the existence of a measure kernel Φ that captures the behaviour of \mathbf{P} in two respects. On the one hand, it describes the evolution in infinitesimal time by fixing a cell splitting rule. On the other hand, it specifies the spatial interdependence of the cells during their overall evolution. The latter gives rise to a Gibbsian description of \mathbf{P} , in that Φ determines the conditional distribution under \mathbf{P} of the evolution of all cells within any bounded window when the evolution of all other cells is given.
- We then turn to a kind of thermodynamic formalism for BRTs \mathbf{P} that are invariant under spatial translations. The basic quantity is an inner entropy density $h^{\text{in}}(\mathbf{P})$, which is defined as the limit of a conditional entropy per unit volume of \mathbf{P} relative to a reference STIT model. The adjective ‘inner’ refers to the fact that only the cells completely inside the respective window are taken into account, rather than all cells that hit the window. The functional h^{in} will be shown to share some familiar

properties of the entropy densities for the standard models of statistical mechanics, at least with some natural adaptations.

- Finally, we consider an arbitrary division kernel Ψ that satisfies some mild regularity conditions, which roughly require that Ψ is not too far from a STIT kernel; such a Ψ will be called moderate. We introduce an associated inner energy density $u^{\text{in}}(\mathbf{P}; \Psi)$ as well as some sort of pressure $v^{\text{in}}(\mathbf{P}; \Psi)$. The resulting inner excess free energy density $h^{\text{in}}(\mathbf{P}; \Psi)$ gives rise to a variational principle, which states that the minimisers of $h^{\text{in}}(\cdot; \Psi)$ are precisely the translation invariant BRTs that admit Ψ as their division kernel. It is further shown that such minimisers do exist, for any prescribed distribution P of the time-zero tessellation. This proves the existence of a BRT \mathbf{P} for any given initial distribution P and any moderate division kernel Ψ . For general Ψ , such a \mathbf{P} is not necessarily unique.

The paper is organised as follows: Section 2 introduces the setup and recalls some necessary facts. Besides tessellations and BRTs, the main concepts are division kernels and local conditional BRTs of Gibbsian type. This section also includes some examples of division kernels to which our theory applies. The main results together with their framework are stated in Section 3. These are Theorem 3.1 on the existence and significance of global division kernels, Theorems 3.5 and 3.6 on the existence of the inner entropy density and its properties, and Theorems 3.9 and 3.10 on the variational characterisation and the existence of invariant BRTs with given moderate division kernels. All proofs are collected in the final Section 4.

2 Preliminaries

2.1 Tessellations

2.1.1 Polytopes and tessellations. Consider the Euclidean space \mathbb{R}^d of arbitrary dimension $d \geq 1$. We shall deal with certain random processes of coloured tessellations of \mathbb{R}^d into (coloured) convex polytopes. Let us specify these terms. First, a *polytope* p in \mathbb{R}^d is the closed convex hull of a finite set of points and is always assumed to have non-empty interior; the set of all such polytopes is denoted by \mathbb{P} . Each polytope $p \in \mathbb{P}$ is equipped with a translation covariant selector $m(p)$, called its ‘*centre*’ or ‘*midpoint*’, for example its barycentre, its Steiner point or its circumcentre. We write $r(p) = \max_{x \in p} |x - m(p)|$ for its *radius* and ∂p and $\text{int}(p)$ for its topological *boundary* resp. *interior*.

More generally, we will assume that each polytope is marked with some internal property, called its *colour*. So, we fix an arbitrary Polish space Σ , which we call the *colour space*. A coloured polytope, called *cell* in the sequel, is a pair $c = (p, \sigma)$ with $p \in \mathbb{P}$ and $\sigma \in \Sigma$. Let us denote by $\text{sp}(c) := p$ and $\text{col}(c) := \sigma$, respectively, the spatial part and the colour of c . The space of cells is thus $\mathbb{C} := \mathbb{P} \times \Sigma$. To simplify notation *we adopt the general convention that spatial operations on cells (and also on coloured tessellations defined below), such as intersections with subsets of \mathbb{R}^d and translations, solely refer to the spatial part and do not affect their colours.* For example, $m(c) := m(\text{sp}(c))$, $r(c) := r(\text{sp}(c))$, $\text{int}(c) := \text{int}(\text{sp}(c))$, $c \cap W := (\text{sp}(c) \cap W, \text{col}(c))$ for $W \subset \mathbb{R}^d$, and $c - x := (\text{sp}(c) - x, \text{col}(c))$ when $x \in \mathbb{R}^d$. Finally, $\text{vol}(c) := \text{vol}_d(\text{sp}(c))$ is the (d -dimensional) volume of the spatial part of c . Let us also define the space $\mathbb{C}_0 = \{c \in \mathbb{C} : m(c) = 0\}$ of cells having their midpoint at the origin.

The cells are the constituents of the coloured tessellations which we introduce now; for brevity we will omit the adjective ‘coloured’ in the following. (Note that letting Σ be a singleton one recovers the uncoloured case usually considered in the literature; cf. [17, 24].)

Definition 2.1. A (coloured) *tessellation* T of \mathbb{R}^d is a countable subset of \mathbb{C} such that

- T is locally finite, in that any bounded subset of \mathbb{R}^d only hits a finite number of cells from T ,
- two distinct cells of T have disjoint interiors, i.e. $\text{int}(c) \cap \text{int}(c') = \emptyset$ for all $c, c' \in T$ with $c \neq c'$,
- the cells cover the whole space, which is to say that $\bigcup_{c \in T} c = \mathbb{R}^d$.

The space of all tessellations of \mathbb{R}^d will henceforth be denoted by \mathbb{T} .

Besides tessellations of \mathbb{R}^d , we will also consider tessellations in local windows $W \subset \mathbb{R}^d$, which will generally be chosen to be polytopes, or sometimes also a finite unions of polytopes. So, we write \mathbb{P}_\cup for the set of all finite, not necessarily connected unions of polytopes, and for $W \in \mathbb{P}_\cup$ we let \mathbb{C}_W be the set of cells that are contained in W . We finally write \mathbb{T}_W for the set of all tessellations of W , that is, of all finite collections $\{c_1, \dots, c_n\}$ of cells with pairwise disjoint interiors and such that $c_1 \cup \dots \cup c_n = W$.

2.1.2 Measurability. We need measurable structures on all spaces introduced above. We start with the space \mathbb{P} of polytopes. As the sets in \mathbb{P} are compact and non-empty, the natural metric on \mathbb{P} is the usual Hausdorff distance d_H ; cf. [17, Chapter 12.3]. Hence, the space \mathbb{P} can be equipped with the Borel σ -field $\mathcal{B}(\mathbb{P})$ induced by d_H . In fact, $\mathcal{B}(\mathbb{P})$ is generated by the sets $\{p \in \mathbb{P} : p \cap B \neq \emptyset\}$ with $B \in \mathcal{B}(\mathbb{R}^d)$, the Borel σ -field on \mathbb{R}^d ; see [17, Chapters 12.2–12.3]. The coloured counterpart \mathbb{C} is endowed with the product σ -field $\mathcal{B}(\mathbb{C}) = \mathcal{B}(\mathbb{P}) \otimes \mathcal{B}(\Sigma)$, where $\mathcal{B}(\Sigma)$ is the Borel σ -field on Σ . The space \mathbb{C}_0 of centred cells receives the trace σ -field.

We next need to introduce a suitable σ -field on \mathbb{T} . As is usual in point process theory, we let $\mathcal{B}(\mathbb{T})$ be the σ -field generated by the counting variables

$$(2.1) \quad N_A : \mathbb{T} \rightarrow \mathbb{N} \cup \{+\infty\}, \quad T \mapsto |T \cap A|, \quad A \in \mathcal{B}(\mathbb{C}),$$

where $|\cdot|$ stands for the cardinality of the argument set, i.e. N_A counts how many cells of T belong to A . In view of the structure of $\mathcal{B}(\mathbb{C})$, $\mathcal{B}(\mathbb{T})$ is also generated by the random variables

$$N_{B,S} : \mathbb{T} \ni T \mapsto |\{c \in T : c \cap B \neq \emptyset, \text{col}(c) \in S\}|,$$

with B a bounded Borel set in \mathbb{R}^d and $S \in \mathcal{B}(\Sigma)$. Moreover, $\mathcal{B}(\mathbb{T})$ is the Borel σ -field for the vague topology on \mathbb{T} , which is generated by the functions

$$e_g : \mathbb{T} \rightarrow [0, \infty), \quad T \mapsto \sum_{c \in T} g(c),$$

where $g \geq 0$ is a continuous function on \mathbb{C} with a bounded support in the spatial coordinate; see [10, Appendix 15.7] or [11, Theorem A2.3].

Definition 2.2. A probability measure P on $(\mathbb{T}, \mathcal{B}(\mathbb{T}))$ satisfying the first-moment condition $\int P(dT) |T_W| < \infty$ for all windows $W \in \mathbb{P}$ is called a *random tessellation*. The set of all such P is denoted by $\mathcal{P}(\mathbb{T})$.

To deal with local properties of tessellations we will often restrict a tessellation to a local window $W \in \mathbb{P}$. We thus define the *projection* to such a W by

$$(2.2) \quad \pi_W : \mathbb{T} \rightarrow \mathbb{T}_W, \quad T \mapsto T_W := \{c \cap W : c \in T, \text{int}(c \cap W) \neq \emptyset\}.$$

In the same manner as above, we may introduce a σ -field $\mathcal{B}(\mathbb{T}_W)$ on \mathbb{T}_W . One can then easily check that the mapping π_W is measurable.

2.2 Branching tessellations

2.2.1 Cutting cells by hyperplanes. We now turn to the main objects of our investigation: tessellations which arise from a given initial tessellation by a successive splitting of cells into two pieces by means of suitable hyperplanes. Recall that a hyperplane η with unit normal $u \in \mathbb{S}_+^{d-1}$ (upper unit half-sphere) and signed distance $r \in \mathbb{R}$ to the origin can be written in the form $\eta = \{x \in \mathbb{R}^d : \langle x, u \rangle = r\}$, where $\langle \cdot, \cdot \rangle$ stands for the usual scalar product. So, the space of hyperplanes can be identified with $\mathbb{S}_+^{d-1} \times \mathbb{R}$. For η as above, we write $\eta^+ = \{x \in \mathbb{R}^d : \langle x, u \rangle \geq r\}$ and $\eta^- = \{x \in \mathbb{R}^d : \langle x, u \rangle \leq r\}$ for the associated half-spaces. More generally, we consider *bicoloured hyperplanes* $H = (\eta, \sigma^+, \sigma^-) \in \mathbb{H} := \mathbb{S}_+^{d-1} \times \mathbb{R} \times \Sigma^2$, for which each of the half-spaces η^\pm is equipped with a colour σ^\pm . We write $\text{sp}(H) := \eta$ and $\text{col}^\pm(H) := \sigma^\pm$, respectively, for the spatial part and the colours of H and again *adopt the convention that spatial operations with bicoloured hyperplanes only refer to the spatial part*, for example, $c \cap H := \text{sp}(c) \cap \text{sp}(H)$ or $c \cap H^\pm := (\text{sp}(c) \cap \text{sp}(H)^\pm, \text{col}^\pm(H))$ for any $c \in \mathbb{C}$. Moreover, for such a cell c , we let

$$(2.3) \quad \langle c \rangle = \{H \in \mathbb{H} : H \cap \text{int}(c) \neq \emptyset\}$$

be the set of all bicoloured hyperplanes, the spatial part of which hits the interior of the spatial part of c . Each bicoloured hyperplane H defines a cell division operation \oslash on tessellations. Namely, let $T \in \mathbb{T}$, $c \in T$ and $H \in \langle c \rangle$. Then \oslash is defined by

$$(2.4) \quad \oslash_{c,H}(T) := (T \setminus \{c\}) \cup \{c \cap H^+, c \cap H^-\}$$

with $c \cap H^\pm$ as above. Branching tessellations are now defined as follows. For simplicity, the time interval will mostly be the unit interval $[0, 1]$.

Definition 2.3. (a) Let $W \in \mathbb{P}_\cup$ be a finite union of polytopes. A *branching tessellation in the window W* with bounded time interval $I = [a, b]$ or $[a, b]$ is a family $\mathbf{T} = (T_s)_{s \in I}$ of tessellations in W such that

- the function $s \mapsto T_s$ from I to \mathbb{T}_W is piecewise constant, right-continuous and has only a finite number of jumps,
- at each point s of discontinuity (so that $T_s \neq T_{s-} := \lim_{r \uparrow s} T_r$), there exists a unique cell $c \in T_{s-}$ and a bicoloured hyperplane $H \in \langle c \rangle$ such that

$$T_s = \oslash_{c,H}(T_{s-}).$$

Further, T_a is called the initial tessellation. We write \mathbb{BT}_W for the set of all such branching tessellations in W .

(b) A family $\mathbf{T} = (T_s)_{0 \leq s \leq 1}$ is called a *branching tessellation in \mathbb{R}^d* if for each window $W \in \mathbb{P}$ the restricted process $\mathbf{T}_W = \pi_W(\mathbf{T}) := (\pi_W(T_s))_{0 \leq s \leq 1}$ is a branching tessellation in W . Again, T_0 is then called the initial tessellation of \mathbf{T} . The set of all branching tessellations in \mathbb{R}^d is denoted by \mathbb{BT} .

The following remark provides a further way of describing the time evolution of a branching tessellation.

Remark 2.4. (a) Let \mathbf{T} be a branching tessellation in a window $W \in \mathbb{P}_\cup$ with time interval $I = [0, 1]$. (The case of other time intervals is similar.) Keeping record of all jump times of \mathbf{T} together with the associated cells that are divided and the respective cutting hyperplanes, one arrives at the set

$$(2.5) \quad \mathcal{D}(\mathbf{T}) = \left\{ (s, c, H) \in (0, 1] \times \mathbb{C} \times \mathbb{H} : T_{s-} \neq T_s, \right. \\ \left. c \in T_{s-}, H \in \langle c \rangle, T_s = \oslash_{c,H}(T_{s-}) \right\}$$




Figure 1: Representation of a two-coloured branching tessellation in a finite window with initial tessellation $T_0 = \{c, c'\}$. The cells living at a time s constitute a tessellation T_s . At each time s_i , a cell that lives up to this moment is selected and cut in two by a bicoloured hyperplane H_i , which impresses its colours onto the cell's pieces.

of all ‘division events’. There is a one-to-one correspondence between \mathbf{T} and the pair $(T_0, \mathcal{D}(\mathbf{T}))$, in that \mathbf{T} can be recovered from the initial tessellation T_0 and the set $\mathcal{D}(\mathbf{T})$ of division events. Indeed, labelling the elements of $\mathcal{D}(\mathbf{T})$ with the indices $1, \dots, n := |\mathcal{D}(\mathbf{T})|$ according to the order of their time coordinates so that $0 =: s_0 < s_1 < \dots < s_n \leq s_{n+1} := 1$, one has the recursion $T_s = T_0$ for $s \in [0, s_1)$ and

$$T_s = \mathcal{O}_{c_i, H_i}(T_{s_{i-1}}) \quad \text{for } s \in [s_i, s_{i+1}), i = 1, \dots, n.$$

Finally, $T_1 = T_{s_n}$.

This description also gives rise to a convenient way of visualising \mathbf{T} as a graph in $[0, 1] \times \mathbb{C}_W$; see Figure 1. The set of vertices is

$$\mathcal{V}(\mathbf{T}) = \{(0, c) : c \in T_0\} \cup \{(s, c \cap H^\pm) : (s, c, H) \in \mathcal{D}(\mathbf{T})\}.$$

Moreover, each $(s, c) \in \mathcal{V}(\mathbf{T})$ is equipped with a ‘lifeline’ $[s, s^*) \times \{c\}$, where $s^* = s'$ if $(s', c, H) \in \mathcal{D}(\mathbf{T})$ for some $s' > s$ and $H \in \langle c \rangle$, and $s^* = 1$ otherwise. If $s^* < 1$, this lifeline is augmented by the lines from (s^*, c) to the two children $(s^*, c \cap H^\pm)$ of (s, c) . If $s^* = 1$, the half-open line $[s, 1) \times \{c\}$ is replaced by the closed line $[s, 1] \times \{c\}$. In this way one obtains a finite forest of binary ‘family’ trees in \mathbb{C}_W that evolve from the cells of T_0 . So, these cells are the roots, or ancestors, and the $|T_0| + |\mathcal{D}(\mathbf{T})|$ leaves form the tessellation T_1 . This branching mechanism is strongly reminiscent of the fragmentation processes considered in [4].

(b) Branching tessellations in the whole space \mathbb{R}^d admit a similar description in terms of division events. For each $\mathbf{T} \in \mathbb{BT}$ we can then define

$$(2.6) \quad \mathcal{D}(\mathbf{T}) = \bigcup_{V \in \mathbb{P}} \bigcap_{W \in \mathbb{P}: W \supset V} \mathcal{D}(\mathbf{T}_W).$$

Conversely, for each $W \in \mathbb{P}$ one can recover the division events in W from $\mathcal{D}(\mathbf{T})$ via

$$\mathcal{D}(\mathbf{T}_W) = \left\{ (s, c \cap W, H) : (s, c, H) \in \mathcal{D}(\mathbf{T}), H \in \langle c \cap W \rangle \right\}.$$

It follows that \mathbf{T} is uniquely determined by T_0 and $\mathcal{D}(\mathbf{T})$, and \mathbf{T} can be regarded as a forest of infinitely many finite binary family trees of coloured cells, the roots of which correspond to the cells of the initial tessellation T_0 of \mathbb{R}^d .

Later on, it will be essential for us to keep track of the past of a branching tessellation. So, instead of considering the evolution $\mathbf{T} = (T_s)_{0 \leq s \leq 1}$ in \mathbb{T} , we will consider the process $(\mathbf{T}_s)_{0 \leq s \leq 1}$ in \mathbb{BT} , which is given by $\mathbf{T}_s = (T_u)_{0 \leq u \leq s}$. Equivalently, \mathbf{T}_s can be thought of as being obtained from \mathbf{T} by removing from $\mathcal{D}(\mathbf{T})$ all elements with time-coordinate larger than s . In this way, each \mathbf{T}_s can be considered to be an element of \mathbb{BT} , which is frozen at time s (and thus remains constant thereafter). The set of all such branching tessellations is denoted by \mathbb{BT}_s . In particular, $\mathbb{BT}_1 = \mathbb{BT}$, and $\mathbb{BT}_u \subset \mathbb{BT}_s$ when $u < s$. We write

$$(2.7) \quad \pi_s : \mathbb{BT} \rightarrow \mathbb{BT}_s, \quad \mathbf{T} \mapsto \mathbf{T}_s,$$

for the natural projection that removes the division events after time s . As before, the non-bold T_s stands for the tessellation at time s , whereas a bold \mathbf{T}_s stands for an element of \mathbb{BT}_s .

Besides this projection concerning time, we have also the projection to a spatial window $W \in \mathbb{P}$, which is given by

$$(2.8) \quad \pi_W : \mathbb{BT} \rightarrow \mathbb{BT}_W, \quad \mathbf{T} \mapsto \mathbf{T}_W = (T_{W,s})_{0 \leq s \leq 1} \quad \text{with} \quad T_{W,s} = \pi_W(T_s),$$

where $\mathbb{BT}_W = \pi_W(\mathbb{BT})$ and π_W is as in (2.2). We also write $\pi_{W,s} = \pi_W \circ \pi_s$, $\mathbf{T}_{W,s} = \pi_{W,s}(\mathbf{T})$ and $\mathbb{BT}_{W,s} = \pi_{W,s}(\mathbb{BT})$. So, to obtain $\mathbf{T}_{W,s}$ from \mathbf{T} one has to remove from $\mathcal{D}(\mathbf{T})$ all division events with a time coordinate exceeding s or a hyperplane not hitting the cell's intersection with W .

2.2.2 Branching random tessellations. Our main objects of interest are probability measures on \mathbb{BT} . So we need to equip \mathbb{BT} with a σ -field. We know from Remark 2.4 that each $\mathbf{T} \in \mathbb{BT}$ is uniquely determined by its initial tessellation T_0 together with the set $\mathcal{D}(\mathbf{T})$ of division events as given by (2.5) and (2.6). Since $\mathcal{D}(\mathbf{T})$ is a locally finite subset of $(0, 1] \times \mathbb{C} \times \mathbb{H}$, one can proceed as usually in point process theory by defining $\mathcal{B} = \mathcal{B}(\mathbb{BT})$ as the smallest σ -field for which the counting variables

$$(2.9) \quad \mathbf{N}_{A,B} : \mathbf{T} \mapsto |T_0 \cap A| + |\mathcal{D}(\mathbf{T}) \cap B|$$

with $A \in \mathcal{B}(\mathbb{C})$ and $B \in \mathcal{B}((0, 1]) \otimes \mathcal{B}(\mathbb{C}) \otimes \mathcal{B}(\mathbb{H})$ are measurable; here $\mathcal{B}((0, 1])$ denotes the Borel σ -field on $(0, 1]$. By standard theory, $(\mathbb{BT}, \mathcal{B})$ is a Borel space. For any window $W \in \mathbb{P}$, we define a σ -field $\mathcal{B}_W = \mathcal{B}(\mathbb{BT}_W)$ on \mathbb{BT}_W in the same way. To simplify notation we will not distinguish between the σ -field \mathcal{B}_W on \mathbb{BT}_W and its preimage $\pi_W^{-1}\mathcal{B}_W$ on \mathbb{BT} , which will be denoted by the same symbol. Anyway, with these definitions it is clear that both the projection π_W in (2.8) and the time restriction map $\pi_\bullet : (s, \mathbf{T}) \mapsto \mathbf{T}_s$ of (2.7) are measurable.

Definition 2.5. A *branching random tessellation (BRT)* of \mathbb{R}^d is a probability measure \mathbf{P} on $(\mathbb{BT}, \mathcal{B})$ satisfying the first-moment condition

$$(2.10) \quad \int \mathbf{P}(d\mathbf{T}) |T_{W,1}| < \infty \quad \text{for all windows } W \in \mathbb{P}.$$

The set of all such BRTs of \mathbb{R}^d is denoted by $\mathcal{P} = \mathcal{P}(\mathbb{BT})$. BRTs within a window $W \in \mathbb{P}_\cup$ are defined analogously.

For every $\mathbf{P} \in \mathcal{P}$ and any of the projections π_* in (2.7) and (2.8) we write $\mathbf{P}_* = \mathbf{P} \circ \pi_*^{-1}$ for the image of \mathbf{P} under π_* . In particular, each \mathbf{P}_s is a BRT. In fact, one can achieve that \mathbf{P}_s depends measurably on s , in that the mapping $[0, 1] \times \mathcal{B} \ni (s, A) \mapsto \mathbf{P}_s(A)$ is a probability kernel, as will be assumed throughout the following. This can be seen by disintegrating the measure

$$(2.11) \quad \overline{\mathbf{P}} := \int_0^1 ds \int \mathbf{P}(d\mathbf{T}) \delta_{(s, \mathbf{T}_s)}$$

on $\overline{\mathbb{BT}} := \{(s, \mathbf{T}_s) : s \in [0, 1], \mathbf{T}_s \in \mathbb{BT}_s\}$; cf. [10, Appendix 15.3]. Later on, we will also consider the projections $\overline{\pi}_W = \text{id} \otimes \pi_W$ that act on the second coordinate of $\overline{\mathbb{BT}}$ as in (2.8) and leave the first coordinate untouched, and the projection images $\overline{\mathbf{P}}_W = \overline{\mathbf{P}} \circ \overline{\pi}_W^{-1}$, where $W \in \mathbb{P}$. We also introduce the notation $\overline{\mathbb{BT}}_W := \overline{\pi}_W(\overline{\mathbb{BT}})$.

2.3 Division kernels

Consider a random element \mathbf{T} of \mathbb{BT}_W for a window $W \in \mathbb{P}_\cup$. The process $(\mathbf{T}_s)_{0 \leq s \leq 1}$ is then automatically Markovian because its ‘past’ is part of the ‘present’. In this paper we will focus on the ‘nice’ case in which the evolution of this Markov process is described by a rate kernel that specifies the jump times and transitions of $(\mathbf{T}_s)_{0 \leq s \leq 1}$. Since the only transitions are single-cell divisions by bicoloured hyperplanes, this means that the rate kernels take the following form.

Definition 2.6. A *division kernel* is a measure kernel Φ from the set

$$\{(s, \mathbf{T}_s, c) \in \overline{\mathbb{BT}} \times \mathbb{C} : c \in T_s\}$$

to \mathbb{H} such that each $\Phi(s, \mathbf{T}_s, c, \cdot)$ is a finite measure supported on $\langle c \rangle \subset \mathbb{H}$. If Φ is only defined for arguments in $\overline{\mathbb{BT}}_W \times \mathbb{C}_W$, Φ is called a *division kernel for the window* $W \in \mathbb{P}_\cup$.

In the following, it will be convenient to work also with the *cumulative division kernel*

$$(2.12) \quad \hat{\Phi}(s, \mathbf{T}_s, \cdot) = \sum_{c \in T_s} \delta_c \otimes \Phi(s, \mathbf{T}_s, c, \cdot)$$

from $\overline{\mathbb{BT}}$ to $\mathbb{C} \times \mathbb{H}$. Note that, conversely, $\Phi(s, \mathbf{T}_s, c, \cdot) = \hat{\Phi}(s, \mathbf{T}_s, \{c\} \times \cdot)$.

The next remark describes how a division kernel determines the evolution of a BRT within a bounded window.

Remark 2.7. *Local BRTs with prescribed division kernels.* Let $W \in \mathbb{P}_\cup$ be a fixed window, Φ_W be a division kernel for W , and $\hat{\phi}_W(s, \mathbf{T}_s) := \hat{\Phi}_W(s, \mathbf{T}_s, T_s \times \langle W \rangle)$ the finite total mass of the cumulative kernel $\hat{\Phi}_W(s, \mathbf{T}_s, \cdot)$. We construct a random element \mathbf{T} of \mathbb{BT}_W as follows:

- (I) Pick an initial tessellation $T_0 \in \mathbb{T}_W$ according to some probability law P_W on \mathbb{T}_W , and let $s_0 = 0$ and $\mathbf{T}_0 := T_0$. Also, let $i = 1$ and proceed with the following random recursion over the number i .
- (R) Suppose that $i \geq 1$ and both a random time $s_{i-1} \in [0, 1]$ and a BRT $\mathbf{T}_{s_{i-1}} \in \mathbb{BT}_{s_{i-1}}$ are already realised. Then take a random time $s_i \in (s_{i-1}, \infty]$ with ‘survival’ probability

$$(2.13) \quad \text{Prob}(s_i > s) = \exp \left[- \int_{s_{i-1}}^s \hat{\phi}_W(u \wedge 1, \mathbf{T}_{s_{i-1}}) du \right] \quad \text{for } s > s_{i-1}.$$

If $s_i \leq 1$, proceed to define an extension $\mathbf{T}_{s_i} \in \mathbb{BT}_{s_i}$ of $\mathbf{T}_{s_{i-1}}$ as follows: Pick a random cell $c_i \in T_{s_{i-1}}$ and a bicoloured hyperplane H_i according to the law

$$\hat{\Phi}_W(s_i, \mathbf{T}_{s_{i-1}}, \cdot) / \hat{\phi}_W(s_i, \mathbf{T}_{s_{i-1}}).$$

(Note that the denominator does not vanish for each possible choice of s_i .) Then let $T_s = T_{s_{i-1}}$ for $s \in (s_{i-1}, s_i)$ and $T_{s_i} = \mathcal{O}_{c_i, H_i}(T_{s_{i-1}})$, that is,

$$\mathcal{D}(\mathbf{T}_{s_i}) = \mathcal{D}(\mathbf{T}_{s_{i-1}}) \cup \{(s_i, c_i, H_i)\}.$$

Next let $i := i+1$ and goto (R). In the case $s_i > 1$ let $T_s = T_{s_{i-1}}$ for $s \in (s_{i-1}, 1]$, set $n = i-1$, and stop.

One needs to ensure that this algorithm terminates after finitely many steps. It is not difficult to show that this is the case if

$$(2.14) \quad \sup_{s, \mathbf{T}_s, c} \Phi_W(s, \mathbf{T}_s, c, \langle c \rangle) =: \phi < \infty;$$

see the proof of Lemma 4.3 below. This lemma shows further that the process $(\mathbf{T}_s)_{0 \leq s \leq 1}$ can be characterised as the unique, in general time-inhomogeneous pure jump (i.e. piecewise constant) Markov process in \mathbb{BT}_W with initial distribution P_W and generator

$$(2.15) \quad \mathbb{L}_{W,s}^{\Phi_W} g(\mathbf{T}_s) = \int_{T_s \times \langle W \rangle} \hat{\Phi}_W(s, \mathbf{T}_s, d(c, H)) \left[g(\mathcal{O}_{s,c,H}(\mathbf{T}_s)) - g(\mathbf{T}_s) \right]$$

at time $s \in [0, 1]$. Here, $\mathcal{O}_{s,c,H}(\mathbf{T}_s) \in \mathbb{BT}_s$

is the branching tessellation that coincides with \mathbf{T}_s for times less than s and equals $\mathcal{O}_{c,H}(T_s)$ at time s , and g is any bounded measurable function on \mathbb{BT}_W . The distribution of \mathbf{T} is a BRT \mathbf{P}_W in $W \in \mathbb{P}$, and this \mathbf{P}_W is called the *BRT in W with division kernel Φ_W and initial distribution P_W* .

2.4 Examples of division kernels

This section contains a few examples of division kernels; two simulation pictures are shown in Figure 2. The first is (by now) classical and will be used as a reference model throughout the following.

Example 2.8. STIT tessellations. Let Λ be a locally finite measure on the set \mathbb{H} of all bicoloured hyperplanes, which is invariant under all translations. That is, under the identification of $H \in \mathbb{H}$ with $(u, r, \sigma^+, \sigma^-) \in \mathbb{S}_+^{d-1} \times \mathbb{R} \times \Sigma^2$, Λ can be written in the form

$$(2.16) \quad \Lambda(dH) = \lambda(du) dr \mu(u, d\sigma^+, d\sigma^-).$$

Here, λ is a measure on \mathbb{S}_+^{d-1} , and μ is a probability kernel from \mathbb{S}_+^{d-1} to Σ^2 . (The translation invariance is expressed by the fact that the r -marginal is Lebesgue measure and μ does not depend on r .) A natural choice is the motion invariant measure Λ_{iso} for which λ is the normalised surface measure λ_{iso} on \mathbb{S}_+^{d-1} and $\mu(u, \cdot) = \nu \otimes \nu$ for a reference probability measure ν on Σ . Then, a STIT tessellation with driving measure Λ is a BRT for the division kernel

$$(2.17) \quad \Lambda^*(s, \mathbf{T}_s, c, \cdot) = \Lambda(\cdot \cap \langle c \rangle).$$

In the uncoloured case, this model has been introduced by Mecke, Nagel and Weiß [12, 13]. Since Λ^* does not depend on the time s , the random holding times $s_i - s_{i-1}$ in Remark 2.7 above are exponentially distributed and can be understood as minima over $c \in T_{s_{i-1}}$ of independent exponential times with parameter $\Lambda(\langle c \rangle)$, which are associated to the presently existing cells. (In the isotropic case $\Lambda = \Lambda_{\text{iso}}$, the parameter $\Lambda(\langle c \rangle)$ is precisely the mean width of c .) In other words, the tessellations evolve according to a continuous-time branching process on \mathbb{C}_W , $W \in \mathbb{P}$, in which all cells c behave independently of each other, live for an exponential time with parameter $\Lambda(\langle c \rangle)$ and then split into two parts according to the conditional distribution $\Lambda(\cdot | \langle c \rangle)$. In particular, this implies that smaller cells live stochastically longer.

In view of this independence of the evolution in different cells, it is clear that for each $T_0 \in \mathbb{T}$ there exists a unique whole-space BRT $\mathbf{\Pi}^\Lambda(T_0, \cdot)$, called STIT tessellation of \mathbb{R}^d , with initial tessellation T_0 and driving measure Λ . In fact, if the support of λ contains a linear basis of \mathbb{R}^d , one can also construct a unique BRT $\mathbf{\Pi}^{\Lambda, \infty} = \mathbf{\Pi}^\Lambda(T_0^\infty, \cdot)$

for the degenerate initial tessellations T_0^∞ that consist of the single ‘cell’ \mathbb{R}^d with any colour σ ; see [12, Theorem 1] and [13, Theorem 1]. (By (2.17), $\mathbf{\Pi}^{\Lambda, \infty}$ does not depend on σ .)

Formally, $\mathbf{\Pi}^\Lambda$ is a probability kernel from \mathbb{T} to \mathbb{BT} . So, for each $P \in \mathcal{P}(\mathbb{T})$, $P\mathbf{\Pi}^\Lambda = \int P(dT_0) \mathbf{\Pi}^\Lambda(T_0, \cdot)$ is the unique BRT for Λ with initial distribution P . Its projections to arbitrary windows $W \in \mathbb{P}$ are given by

$$(2.18) \quad (P\mathbf{\Pi}^\Lambda) \circ \pi_W^{-1} = P_W \mathbf{\Pi}_W^\Lambda$$

for the restricted STIT kernel $\mathbf{\Pi}_W^\Lambda(T_{W,0}, \cdot)$ from \mathbb{T}_W to \mathbb{BT}_W with the restricted driving measure $\Lambda(\cdot \cap \langle W \rangle)$. The abbreviation STIT stands for **st**ability under the operation of **i**teration of tessellations. An explanation and further remarkable properties can be found in [12], [13], [15], [18]–[23] and [25].

A generalisation of the STIT models, which still keeps the independence of the division process for distinct cells, are the *cell-driven BRTs*, which have division kernels of the form

$$(2.19) \quad \Phi(s, \mathbf{T}_s, c, dH) = \varphi(c, H) \Lambda(dH)$$

with a density function $\varphi(c, H)$ on $\mathbb{C} \times \mathbb{H}$ which vanishes except when $H \in \langle c \rangle$. A special case are the *shape-driven BRTs* investigated in [20]; see also the examples therein.

The next example demonstrates the flexibility of modelling in the present setting: It combines an interaction between the colours of the cells with a geometric homogenisation mechanism and an aging effect. The last feature takes advantage of the fact that division kernels may also depend on the past.

Example 2.9. *Contact-induced mutations with size balancing and aging.* Let the colour space be $\Sigma = \{-1, 1\}$ and consider a division kernel of the form

$$\Phi(s, \mathbf{T}_s, c, dH) = \varphi(c, \eta) \lambda_{\text{iso}}(du) dr \mu(s, \mathbf{T}_s, c, d\sigma^+) \mu(s, \mathbf{T}_s, c, d\sigma^-),$$

where $H = (\eta, \sigma^+, \sigma^-)$ with spatial part $\eta = (u, r) \in \mathbb{S}_+^{d-1} \times \mathbb{R}$ and colours $\sigma^\pm \in \Sigma$. A special choice of the geometric prefactor is

$$\varphi(c, \eta) = \varepsilon \mathbb{1}_{\langle c \rangle}(\eta) + \varepsilon^{-1} \mathbb{1}_{\langle \varepsilon \star c \rangle}(\eta),$$




Figure 2: Simulations of two BRTs with isotropic selection of lines, two colours and the full window as single initial cell. *Left:* A STIT tessellation; colours are chosen at random. *Right:* Colour mutation and size balancing as in Example 2.9, but without aging. Here, $\varepsilon = 0.05$, $\beta(s) = (1 + s)/2$ and a mixed boundary condition as indicated.

for some small $\varepsilon > 0$; here, $\varepsilon \star c = m(c) + \varepsilon(c - m(c))$ is the ε -retraction of c . This choice has the effect that the cutting hyperplane will typically pass close to the midpoint $m(c)$ of c , so that its two daughter cells have comparable size. One can further choose the colouring rule

$$\mu(s, \mathbf{T}_s, c, \cdot) = \delta_{\text{col}(c)} + \beta(\mathbf{a}_{s,c}, \mathbf{T}_s, \mathbf{s}_{c,T_s}) \delta_{-\text{col}(c)},$$

where $\mathbf{a}_{s,c}, \mathbf{T}_s = s - \min\{u \in [0, s] : c \in T_u\}$ is the age of c at time s ,

$$\mathbf{s}_{c,T_s} = \sum_{c' \in T_s : \text{col}(c') = -\text{col}(c)} \text{vol}_{d-1}(c \cap c') / \text{vol}_{d-1}(\partial c)$$

is the opposite-type surface fraction (measured by the Hausdorff measure of dimension $d-1$), and $\beta : [0, 1]^2 \rightarrow (0, \infty)$ is a suitable positive function. For instance, β can be taken to be decreasing in \mathbf{a} so that increasing age reduces the willingness of splitting and mutating. One can further let β be increasing in \mathbf{s} . Then, the larger a cell's surface fraction is in contact with cells of opposite type, the more the cell gets 'nervous' and hurries to divide, and the more likely it is that its daughter cells mutate to adapt their type to that of the neighbours.

Our third example may seem somewhat exotic. It will be used in Remark 3.12 to demonstrate that a BRT on the full space \mathbb{R}^d is not necessarily uniquely determined by its initial distribution and its division kernel.

Example 2.10. *Directional infinite-range interaction.* This is an uncoloured model, for which Σ is a singleton. We further confine ourselves to the planar case $d = 2$. Let $\Lambda_{\text{hor}}(dH) = \delta_{(0,1)}(du)dr$ and $\Lambda_{\text{vert}}(dH) = \delta_{(1,0)}(du)dr$ be the measures on $\mathbb{H} = \mathbb{S}_+^1 \times \mathbb{R}$ for which all lines are horizontal resp. vertical. For any cell $c \in \mathbb{C}$ let $\text{diam}_{\text{hor}}(c) = \max_{x,y \in c} |x_1 - y_1|$ and $\text{diam}_{\text{vert}}(c) = \max_{x,y \in c} |x_2 - y_2|$ be the horizontal and vertical diameters of c , where x_i stands for the i th coordinate of x . Also, let

$$\mathbb{C}_{\text{hor}} = \{c \in \mathbb{C} : \text{diam}_{\text{hor}}(c) > \text{diam}_{\text{vert}}(c)\}$$

be the set of all 'horizontal' cells. Finally, writing $[n]$ for the centred square of area n^2 , let

$$\rho_{\text{hor}}(T) = \limsup_{n \rightarrow \infty} n^{-2} |\{c \in T \cap \mathbb{C}_{\text{hor}} : m(c) \in [n]\}|$$

be the upper density of horizontal cells for a tessellation $T \in \mathbb{T}$, and define $\rho_{\text{vert}}(T)$ analogously. Then let

$$\mathbb{T}_{\text{hor}} = \{T \in \mathbb{T} : \rho_{\text{hor}}(T) > \rho_{\text{vert}}(T)\}$$

be the set of tessellations with a dominating fraction of horizontal cells, and $\mathbb{T}_{\text{vert}} = \mathbb{T} \setminus \mathbb{T}_{\text{hor}}$. Consider the division kernel

$$(2.20) \quad \Phi(s, \mathbf{T}_s, c, \cdot) = \mathbf{1}_{\mathbb{T}_{\text{hor}}}(T_s) \Lambda_{\text{hor}}(\cdot \cap \langle c \rangle) + \mathbf{1}_{\mathbb{T}_{\text{vert}}}(T_s) \Lambda_{\text{vert}}(\cdot \cap \langle c \rangle).$$

Since \mathbb{T}_{hor} is invariant under translations and tail-measurable, this Φ looks at the actual tessellation 'at infinity' in order to decide whether the cutting line should be horizontal or vertical.

2.5 Conditional BRTs

In this section we introduce a Gibbsian perspective on BRTs. As is standard in the theory of Gibbs measures, one aims at describing a macroscopic system by means of its local conditional distributions that describe the behaviour inside a bounded region

when the remaining system is fixed. Here we define such conditional distributions in the context of BRTs. Let $W \in \mathbb{P}$ be a fixed window.

2.5.1 Inner and outer projections. Recall from (2.2) and (2.8) that the projections π_W and π_W^* are defined by intersecting the cells with W , and thus wipes off much information on the cell geometry (such as, e.g., the location of midpoints). To avoid this, we introduce the ‘inner’ projection

$$(2.21) \quad \pi_W^{\text{in}} : \mathbb{T} \rightarrow \mathbb{T}_W^{\text{in}}, \quad T \mapsto T_W^{\text{in}} := \{c \in T : c \subset \text{int}(W)\},$$

which removes all cells which are not completely contained in the interior of W . It takes values in the set \mathbb{T}_W^{in} of all possibly empty, not necessarily connected collections of cells inside W with pairwise disjoint interiors. The counting variables N_A in (2.1) are even defined on \mathbb{T}_W^{in} and generate a σ -field $\mathcal{B}(\mathbb{T}_W^{\text{in}})$, for which π_W^{in} is measurable. As the cells of T_W^{in} are even required to be contained in the interior of W , T_W^{in} is a measurable function of T_W .

In the same way, we define the *inner projection*

$$(2.22) \quad \pi_W^{\text{in}} : \mathbb{BT} \ni \mathbf{T} \mapsto \mathbf{T}_W^{\text{in}} = (T_{W,s}^{\text{in}})_{0 \leq s \leq 1}$$

on \mathbb{BT} , where $T_{W,s}^{\text{in}} = \pi_W^{\text{in}}(T_s)$. Arguing as in Remark 2.4, one finds that \mathbf{T}_W^{in} is uniquely determined by $T_{W,0}^{\text{in}}$, $\mathcal{D}(\mathbf{T}_W^{\text{in}})$, and the finite set of all ‘immigration events’ (s, c) with $T_{W,s}^{\text{in}} = T_{W,s-}^{\text{in}} \cup \{c\}$. Consequently, one can generate a σ -field on the range $\mathbb{BT}_W^{\text{in}}$ of π_W^{in} by means of counting variables similar to those in (2.9), so that π_W^{in} becomes measurable. Note also that $\pi_W^{\text{in}} = \pi_W^{\text{in}} \circ \pi_W$.

Complementary to the above, we also introduce an ‘outer’ projection for W by

$$(2.23) \quad \pi_W^{\text{out}} : \mathbb{T} \ni T \mapsto T_W^{\text{out}} := T \setminus T_W^{\text{in}} = \{c \in T : c \setminus \text{int}(W) \neq \emptyset\},$$

and a ‘boundary’ projection

$$(2.24) \quad \pi_W^{\partial} : \mathbb{T} \ni T \mapsto \pi_W^{\partial}(T_W) = \{c \cap W : c \in T, c \cap \partial W \neq \emptyset\}.$$

Likewise, on the level of branching tessellations we define

$$(2.25) \quad \pi_W^{\text{out}} : \mathbb{BT} \ni \mathbf{T} \mapsto \mathbf{T}_W^{\text{out}} = (\pi_W^{\text{out}}(T_s))_{0 \leq s \leq 1},$$

$$(2.26) \quad \pi_W^{\partial} : \mathbb{BT} \ni \mathbf{T} \mapsto \mathbf{T}_W^{\partial} = (\pi_W^{\partial}(T_s))_{0 \leq s \leq 1} = \pi_W^{\text{out}}(\mathbf{T}_W).$$

In the forest picture of Figure 1, each $\mathbf{T}_W^{\text{out}}$ in the range $\mathbb{BT}_W^{\text{out}}$ of π_W^{out} corresponds to a forest of binary trees from which all cells within W are erased. So, one can use the counting variables in (2.9) to generate a σ -field on $\mathbb{BT}_W^{\text{out}}$, and π_W^{out} is then evidently measurable. The same applies to π_W^{∂} . Furthermore, to keep the full information on the initial tessellation in \mathbb{R}^d resp. in W , it will also be convenient to introduce the mappings

$$(2.27) \quad \pi_W^{0,\text{out}} : \mathbf{T} \mapsto (T_{W,0}^{\text{in}}, \mathbf{T}_W^{\text{out}}),$$

$$(2.28) \quad \pi_W^{0,\partial} : \mathbf{T} \mapsto (T_{W,0}^{\text{in}}, \mathbf{T}_W^{\partial}).$$

For each of the projections π_W^* in (2.22), (2.25), (2.26), (2.27), and (2.28), we write $\mathcal{B}_W^* = \sigma(\pi_W^*)$ for the σ -field on \mathbb{BT} that is generated by this projection. By abuse of notation, we will use the same symbol \mathcal{B}_W^* for the σ -field on the range of π_W^* .

2.5.2 Conditional BRTs. Let $\mathbf{T} \in \mathbb{BT}$ any branching tessellation. Consider the time-dependent ‘inner’ window

$$(2.29) \quad \text{in}_W(s, \mathbf{T}_W^{\partial}) := W \setminus \text{int}(\bigcup \{c : c \in T_{W,s}^{\partial}\}) = \bigcup \{c : c \in T_{W,s}^{\text{in}}\},$$

which is possibly empty and not necessarily connected. It is measurable jointly in both arguments, piecewise constant and right-continuous as a function of s . Let

$$(2.30) \quad 0 < t_1 = t_1(\mathbf{T}_W^\partial) < \cdots < t_n = t_n(\mathbf{T}_W^\partial) < 1$$

be the jump times of the path $s \mapsto \text{in}_W(s, \mathbf{T}_W^\partial)$, which depend measurably on \mathbf{T}_W^∂ . (Note that possibly $n(\mathbf{T}_W^\partial) = 0$. For the sake of convenience, we also exclude the case that there is a jump at time 1, which occurs with probability zero.) At each t_i , \mathbf{T}_W^∂ creates a new cell c_i inside W , namely

$$c_i = c_i(\mathbf{T}_W^\partial) := \text{cl}(\text{in}_W(t_i, \mathbf{T}_W^\partial) \setminus \text{in}_W(t_{i-1}, \mathbf{T}_W^\partial)),$$

where $t_0 = 0$. In other words, \mathbf{T}_W^∂ induces a process of immigration of cells into W .

Definition 2.11. Let Φ be a division kernel and suppose that the following random process $\mathbf{S} = (S_s)_{0 \leq s \leq 1}$ with $\mathbf{S} \cup \mathbf{T}_W^{\text{out}} := (S_s \cup T_{W,s}^{\text{out}})_{0 \leq s \leq 1} \in \mathbb{BT}$ is well-defined:

- Let $\mathbf{S}_{[0,t_1)} = (S_s)_{0 \leq s < t_1}$ be the BRT in the window $\text{in}_W(0, \mathbf{T}_W^\partial)$ with time interval $[0, t_1)$, initial tessellation $S_0 = T_{W,0}^{\text{in}}$ and division kernel

$$\Phi_W^{\text{in}}(s, \mathbf{S}_s, c, \cdot | \mathbf{T}_W^{\text{out}}) := \Phi(s, \mathbf{S}_s \cup \mathbf{T}_{W,s}^{\text{out}}, c, \cdot) \quad \text{for } c \in S_s, s \in [0, t_1).$$

- For $i = 1, \dots, n$ and conditional on \mathbf{S}_{t_i-} let $\mathbf{S}_{[t_i, t_{i+1})} = (S_s)_{t_i \leq s < t_{i+1}}$ be the BRT in the window $\text{in}_W(t_i, \mathbf{T}_W^\partial) = c_i \cup \text{in}_W(t_{i-1}, \mathbf{T}_W^\partial)$ with time interval $[t_i, t_{i+1})$, initial tessellation $S_{t_i} = S_{t_i-} \cup \{c_i\}$ and division kernel

$$\Phi_W^{\text{in}}(s, \mathbf{S}_s, c, \cdot | \mathbf{T}_W^{\text{out}}) := \Phi(s, \mathbf{S}_s \cup \mathbf{T}_{W,s}^{\text{out}}, c, \cdot) \quad \text{for } c \in S_s, s \in [t_i, t_{i+1}).$$

Here $t_{n+1} = 1$, and we finally set $S_1 := S_{1-}$.

The distribution of \mathbf{S} on $\mathbb{BT}_W^{\text{in}}$ will be denoted by $\mathbf{G}_W^\Phi(\cdot | \mathbf{T}_W^{0,\text{out}})$ and is called the *conditional BRT for Φ in W with initial tessellation $T_{W,0}^{\text{in}}$ and boundary condition $\mathbf{T}_W^{\text{out}}$* .

By construction, \mathbf{G}_W^Φ is a probability kernel from $(\mathbb{BT}_W^{0,\text{out}}, \mathcal{B}_W^{0,\text{out}})$ to $(\mathbb{BT}_W^{\text{in}}, \mathcal{B}_W^{\text{in}})$.

Example 2.12. *Conditional STIT tessellations.* As in Example 2.8, let Λ be a locally finite measure on \mathbb{H} and Λ^* be the associated division kernel; cf. (2.17). Then, $\mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\text{out}}) := \mathbf{G}_W^{\Lambda^*}(\cdot | \mathbf{T}_W^{0,\text{out}})$ is simply the distribution of

$$\bigcup_{c \in T_{W,0}^{\text{in}}} \mathbf{S}^{(c)} \cup \bigcup_{i=1}^n \mathbf{S}^{(i)}$$

for independent random STIT tessellations $\mathbf{S}^{(c)}$ and $\mathbf{S}^{(i)}$ for Λ . Here, $\mathbf{S}^{(c)}$ evolves in time $[0, 1]$ from the single-cell tessellation $S_0^{(c)} = \{c\}$ of the initial polytope $\text{sp}(c)$, whereas $\mathbf{S}^{(i)}$ evolves in time $[t_i, 1]$ from the single-cell initial tessellation $S_{t_i}^{(i)} = \{c_i\}$ of the ‘immigrated’ polytope $\text{sp}(c_i)$ and is extended to the full interval $[0, 1]$ by setting $S_s^{(i)} = \emptyset$ for $s \in [0, t_i)$. Since Λ^* does not depend on the surrounding tessellation, it follows that the measure $\mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\text{out}})$ depends only on $\mathbf{T}_W^{0,\partial}$.

Here is the natural counterpart of the concept of (macroscopic) Gibbs measures in our setup of branching random tessellations.

Definition 2.13. Let Φ be any division kernel. A BRT $\mathbf{P} \in \mathcal{P}$ is called a *Gibbsian BRT for Φ* if $\mathbf{P}_W^{0,\text{out}} \otimes \mathbf{G}_W^\Phi = \mathbf{P}$ for all $W \in \mathbb{P}$ or, more explicitly, if

$$\mathbf{P}(\pi_W^{\text{in}} \in A | \mathcal{B}_W^{0,\text{out}})(\mathbf{T}) = \mathbf{G}_W^\Phi(A | \mathbf{T}_W^{0,\text{out}}) \quad \text{for } \mathbf{P}\text{-almost all } \mathbf{T}$$

for all $A \in \mathcal{B}_W^{\text{in}}$ and $W \in \mathbb{P}$.

2.6 Translation invariance

A main focus of this paper is on BRTs that are invariant under spatial translations. For each $x \in \mathbb{R}^d$, the translation ϑ_x by the vector $-x$ acts

- on cells $c \in \mathbb{C}$ via $\vartheta_x : c \mapsto c - x := (\text{sp}(c) - x, \text{col}(c))$,
- on bicoloured hyperplanes $H \in \mathbb{H}$ via

$$\vartheta_x : H \mapsto H - x := (\text{sp}(H) - x, \text{col}^+(H), \text{col}^-(H)),$$

- on tessellations $T \in \mathbb{T}$ via

$$\vartheta_x : T \mapsto T - x := \{c - x : c \in T\},$$

- on branching tessellations $\mathbf{T} = (T_s)_{0 \leq s \leq 1} \in \mathbb{BT}$ via

$$\vartheta_x : \mathbf{T} \mapsto \mathbf{T} - x := (T_s - x)_{0 \leq s \leq 1}.$$

That is, only the spatial coordinates are shifted, but the colours remain unchanged. Moreover, by abuse of notation we use the same symbol ϑ_x for the translation on each level, and we will also use it for the simultaneous translation of pairs of objects as above.

Definition 2.14. A BRT $\mathbf{P} \in \mathcal{P}$ is called *translation invariant* if it is invariant under the action of the translation group $\Theta = (\vartheta_x)_{x \in \mathbb{R}^d}$ on \mathbb{BT} , in that $\mathbf{P} \circ \vartheta_x^{-1} = \mathbf{P}$ for all $x \in \mathbb{R}^d$. We write $\mathcal{P}_\Theta = \mathcal{P}_\Theta(\mathbb{BT})$ for the set of all translation invariant BRTs that satisfy the first-moment condition (2.10), which by translation invariance is equivalent to the requirement that the ‘hitting intensity’

$$(2.31) \quad i_1(\mathbf{P}) := \int \mathbf{P}(\mathrm{d}\mathbf{T}) |T_{[1],1}|$$

is finite. Here, $[1] := [-1/2, 1/2]^d$ stands for the centred unit cube.

Translation invariance allows to investigate the behaviour of a random tessellation ‘around a typical cell’, which for convenience is located ‘around the origin’. This is formalised by means of Palm calculus as presented in [10, Chapter 12] and [17, Theorem 4.1.1]. Let $\mathbf{P} \in \mathcal{P}_\Theta$ be given. Then, the *Campbell measure* of \mathbf{P} on $\mathbb{BT} \times \mathbb{C}$ is defined by

$$(2.32) \quad \mathbf{C}^{\mathbf{P}} = \int \mathbf{P}(\mathrm{d}\mathbf{T}) \sum_{c \in T_1} \delta_{(\mathbf{T}, c)}.$$

It captures the joint distribution of the (terminal) cells and the complete history of their surrounding tessellation. The Palm calculus now states that there exists a finite measure \mathbf{P}^0 on $\mathbb{BT} \times \mathbb{C}_0$, the so-called *Palm measure* of \mathbf{P} , such that the *Palm formula*

$$(2.33) \quad \int \mathrm{d}\mathbf{C}^{\mathbf{P}}(\mathbf{T}, c) f(m(c), c - m(c), \mathbf{T} - m(c)) = \int \mathrm{d}x \int \mathrm{d}\mathbf{P}^0(\mathbf{T}, c) f(x, c, \mathbf{T})$$

holds for any non-negative measurable function f on $\mathbb{R}^d \times \mathbb{C}_0 \times \mathbb{T}$. Its normalised marginal on \mathbb{C}_0 is called the *typical cell distribution*.

Later on, we will often consider the integral over time s of the Campbell measure and the Palm measure of the projected BRTs \mathbf{P}_s , and it will be convenient to have a shorthand notation for these objects. So we define the *extended Campbell measure*

$$(2.34) \quad \overline{\mathbf{C}}^{\mathbf{P}} = \int_0^1 \mathrm{d}s \int \mathbf{P}_s(\mathrm{d}\mathbf{T}_s) \sum_{c \in T_s} \delta_{(s, \mathbf{T}_s, c)}$$

and the *extended Palm measure*

$$(2.35) \quad \bar{\mathbf{P}}^0 = \int_0^1 ds \int d\mathbf{P}_s^0(\mathbf{T}_s, c) \delta_{(s, \mathbf{T}_s, c)}.$$

For $W \in \mathbb{P}$ we similarly define the *extended local Campbell measure*

$$(2.36) \quad \bar{\mathbf{C}}^{\mathbf{P}^W} = \int_0^1 ds \int \mathbf{P}_{W,s}(d\mathbf{T}_{W,s}) \sum_{c \in T_{W,s}} \delta_{(s, \mathbf{T}_{W,s}, c)}.$$

Also, we will often use the *time-integrated* version of the Palm formula (2.33), where the Campbell measure $\mathbf{C}^{\mathbf{P}}$ and the Palm measure \mathbf{P}^0 are replaced by their extended relatives $\bar{\mathbf{C}}^{\mathbf{P}}$ and $\bar{\mathbf{P}}^0$, respectively. For example, combining the time-integrated Palm formula with the first-moment condition (2.31) we find that the total mass of $\bar{\mathbf{P}}^0$ can be estimated by

$$(2.37) \quad \|\bar{\mathbf{P}}^0\| = \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) |\{c \in T_s : m(c) \in [1]\}| \leq i_1(\mathbf{P}) < \infty$$

We conclude this section with some comments on random, but not branching, tessellations $P \in \mathcal{P}(\mathbb{T})$. These can be considered as BRTs by identifying the space \mathbb{T} with \mathbb{BT}_0 . In particular, it is then clear what translation invariance means, and we can introduce the set $\mathcal{P}_\Theta(\mathbb{T})$ of all translation invariant random tessellations P that satisfy the first-moment condition

$$(2.38) \quad i_0(P) := \int P(dT) |T_{[1]}| < \infty.$$

Since $i_0(\mathbf{P} \circ \pi_0^{-1}) \leq i_1(\mathbf{P})$, the initial distribution of each $\mathbf{P} \in \mathcal{P}_\Theta$ satisfies (2.38).

3 Results

Our results use a STIT tessellation as a reference model. Therefore we fix a locally finite reference measure Λ on \mathbb{H} which is invariant under translations. Moreover, we write $\Pi^\Lambda(T_0, \cdot)$ for the associated STIT kernel, as introduced ibidem.

3.1 Division kernels for BRTs

As we have seen in Remark 2.7, the concept of division kernels allows to introduce a convenient class of local BRTs. Our first goal here is a global counterpart. That is, we ask whether the evolution of a global BRT $\mathbf{P} \in \mathcal{P}$ can be described by a division kernel Φ . As we will see, this is the case whenever \mathbf{P} is **locally absolutely continuous** with respect to the STIT model $P\Pi^\Lambda$ with initial distribution $P = \mathbf{P} \circ \pi_0^{-1}$, in that

$$(LAC) \quad \mathbf{P}_W \ll P_W \Pi_W^\Lambda \quad \text{for all } W \in \mathbb{P};$$

recall that $P_W \Pi_W^\Lambda = (P\Pi^\Lambda) \circ \pi_W^{-1}$ by (2.18).

We note in passing that (LAC) also implies that the realisations of \mathbf{P} almost surely exhibit a ‘tame’ geometry. Namely, in the planar case, they show exactly one type of vertices, the so-called T-vertices, at which an endpoint of a line segment hits an inner point of another line segment (provided this holds already for the initial tessellation); see [13, 15] and the references cited therein.

To state the result we recall that the extended measure $\bar{\mathbf{P}}$ and the extended projections $\bar{\pi}_W$ have been introduced in and after (2.11), and that the operator $\mathcal{O}_{s,c,H}$ is

defined after (2.15). Further, we will need a projection that refers to the cell division procedure. Namely, for $W \in \mathbb{P}$ we introduce the projection

$$(3.1) \quad \tilde{\pi}_W : (c, H) \mapsto (c \cap W, H)$$

on $\mathbb{C} \times \mathbb{H}$, which for each $T \in \mathbb{T}$ maps the set

$$\tilde{\pi}_W^{-1} \Delta(T_W) := \{(c, H) : c \in T, H \in \langle c \cap W \rangle\}$$

onto $\Delta(T_W) := \{(c, H) : c \in T_W, H \in \langle c \rangle\}$. Finally, a real function on \mathbb{BT} is called *local* if it is \mathcal{B}_W -measurable for some $W \in \mathbb{P}$.

Theorem 3.1. *For each $\mathbf{P} \in \mathcal{P}$ satisfying (LAC) there exists a division kernel Φ for \mathbf{P} such that the following holds.*

- (a) *For all non-negative measurable functions f on $\overline{\mathbb{BT}} \times \mathbb{C} \times \mathbb{H}$,*

$$\begin{aligned} \int \mathbf{P}(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T})} f(s, \mathbf{T}_{s-}, c, H) \\ = \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) \int \hat{\Phi}(s, \mathbf{T}_s, d(c, H)) f(s, \mathbf{T}_s, c, H). \end{aligned}$$

- (b) *The (in general time-inhomogeneous) Markov process $\mathbf{T} = (\mathbf{T}_s)_{0 \leq s \leq 1}$ in \mathbb{BT} with distribution \mathbf{P} satisfies the forward equation with generator*

$$\mathbb{L}_s^\Phi g(\mathbf{T}_s) = \int \hat{\Phi}(s, \mathbf{T}_s, d(c, H)) [g(\odot_{s, c, H}(\mathbf{T}_s)) - g(\mathbf{T}_s)],$$

at time $0 \leq s \leq 1$; here, g is any bounded local function on \mathbb{BT} .

- (c) *\mathbf{P} is a Gibbsian BRT for Φ .*
- (d) *For every $W \in \mathbb{P}$, the local projection \mathbf{P}_W is a BRT in W for the localised division kernel $\hat{\Phi}_W$, which is defined as a regular version of the conditional measure*

$$\hat{\Phi}_W(s, \mathbf{T}_{W, s}, B) := \mathbb{E}_{\bar{\mathbf{P}}}[\hat{\Phi}(\cdot, \cdot, \tilde{\pi}_W^{-1} B) \mid \bar{\pi}_W = (s, \mathbf{T}_{W, s})],$$

where $s \in [0, 1]$ and B is any measurable subset of $\mathbb{C}_W \times \langle W \rangle$.

Definition 3.2. If \mathbf{P} and Φ are such that the properties (a) to (d) above hold, we will say that \mathbf{P} admits the (global) division kernel Φ , or that Φ is a division kernel for \mathbf{P} .

Statement (a) characterises the ‘jump intensity measure’ in terms of Φ . It implies statements (b) and (c), which describe the temporal and the spatial aspect of \mathbf{P} , respectively. So, (b) asserts that the global process with distribution \mathbf{P} evolves in the very same way as the local processes in Remark 2.7, in that a cell c in environment \mathbf{T}_s at time s is split by a bicoloured hyperplane H with instantaneous intensity $\hat{\Phi}(s, \mathbf{T}_s, d(c, H))ds$. On the other hand, by the definition of Gibbsian BRTs, statement (c) means that Φ also describes the cell splitting mechanism within a local window when the evolution of all other cells is given. The final statement (d) addresses the (unconditioned) marginal process in a local window W and shows that its division kernel is obtained by the natural averaging over the possible environments outside W . (The converse question of whether a BRT for a given division kernel exists will be addressed in Theorem 3.10.)

We can also state that the division kernel for a translation invariant BRT can be chosen to be covariant.

Corollary 3.3. *Each $\mathbf{P} \in \mathcal{P}_\Theta$ satisfying (LAC) admits a division kernel Φ which is covariant under translations, in the sense that*

$$(3.2) \quad \hat{\Phi}(s, \mathbf{T}_s, \vartheta_x^{-1} \cdot) = \hat{\Phi}(s, \mathbf{T}_s - x, \cdot)$$

for all $x \in \mathbb{R}^d$ and all $(s, \mathbf{T}_s) \in \overline{\mathbb{BT}}$. Moreover, the local division kernels can be chosen to be covariant in the sense that

$$(3.3) \quad \hat{\Phi}_W(s, \mathbf{T}_{W,s}, \vartheta_x^{-1} \cdot) = \hat{\Phi}_{W-x}(s, \mathbf{T}_{W,s} - x, \cdot)$$

for all $x \in \mathbb{R}^d$, $(s, \mathbf{T}_{W,s}) \in \overline{\mathbb{BT}}_W$ and $W \in \mathbb{P}$.

3.2 The inner entropy density

We now turn to a ‘thermodynamic’ investigation of translation invariant BRTs. Our goal here is an appropriate notion of entropy. Recall that the relative entropy, or Kullback-Leibler divergence, between two probability measures μ and ν on a common measurable space is defined to be $\mathcal{H}(\mu; \nu) = \int \log f \, d\mu$ if $\mu \ll \nu$ with Radon-Nikodym density f , and $+\infty$ otherwise. It can also be written in the form

$$(3.4) \quad \mathcal{H}(\mu; \nu) = \int \varrho(f) \, d\nu,$$

where ϱ is the non-negative convex function

$$(3.5) \quad \varrho : a \mapsto 1 - a + a \log a.$$

The formula (3.4) readily shows that $\mathcal{H}(\mu; \nu) \geq 0$ with equality precisely when $\mu = \nu$. We can also take it as the definition of relative entropy in the more general case when μ and ν are finite, not necessarily normalised measures.

Further, if \mathcal{A} is a sub- σ -field of the underlying σ -field then the *conditional relative entropy given \mathcal{A}* is defined as

$$(3.6) \quad \mathcal{H}(\mu; \nu | \mathcal{A}) = \int \mathcal{H}(\mu_{\mathcal{A}}(\cdot | x); \nu_{\mathcal{A}}(\cdot | x)) \, \mu(dx),$$

where $\mu_{\mathcal{A}}(\cdot | x)$ and $\nu_{\mathcal{A}}(\cdot | x)$ are conditional measure kernels given \mathcal{A} for μ and ν , respectively (provided such kernels exist).

In our setup, we take the STIT model for Λ as our reference measure and introduce an ‘inner’ entropy as follows. Recall the definition (2.28) of $\pi_W^{0,\partial}$ and its associated σ -field $\mathcal{B}_W^{0,\partial} = \sigma(\pi_W^{0,\partial})$, and Example 2.12 for the definition of the kernel \mathbf{G}_W^Λ .

Definition 3.4. Let $\mathbf{P} \in \mathcal{P}$ be a BRT and $W \in \mathbb{P}$. The *inner entropy of \mathbf{P} in W* is then defined by

$$(3.7) \quad \mathcal{H}_W^{\text{in}}(\mathbf{P}) := \mathcal{H}(\mathbf{P}_W; P_{W,0} \Pi_W^\Lambda | \mathcal{B}_W^{0,\partial}) = \mathcal{H}(\mathbf{P}_W; \mathbf{P}_W^{0,\partial} \otimes \mathbf{G}_W^\Lambda).$$

(According to physical convention we should add a minus sign, but here we prefer to ignore this convention.)

So, the attribute ‘inner’ means that this entropy compares the evolution of \mathbf{P} with that of the STIT model only for those cells that are completely contained in W , while the evolution of all other cells hitting W is ignored. In the setup of Gibbs measures on general graphs, an analogous concept of inner entropy has been considered by Föllmer and Snell [8].

Next, let $[n] := [-n/2, n/2]^d$ denote the closed centred cube of volume n^d . For a translation invariant BRT $\mathbf{P} \in \mathcal{P}_\Theta$, one expects that the limiting inner entropy per unit volume

$$\lim_{n \rightarrow \infty} n^{-d} \mathcal{H}_{[n]}^{\text{in}}(\mathbf{P})$$

exists, which is then called the *inner entropy density* of \mathbf{P} (relative to the reference STIT cutting rule Λ). Indeed, our result is the following; see (2.35) for the definition of the extended Palm measure $\bar{\mathbf{P}}^0$.

Theorem 3.5. *For each $\mathbf{P} \in \mathcal{P}_\Theta$ there exists the possibly infinite limit*

$$h^{\text{in}}(\mathbf{P}) := \lim_{n \rightarrow \infty} n^{-d} \mathcal{H}_{[n]}^{\text{in}}(\mathbf{P}).$$

If this limit is finite, \mathbf{P} admits a translation covariant division kernel Φ , and

$$\begin{aligned} (3.8) \quad h^{\text{in}}(\mathbf{P}) &= \mathcal{H}(\bar{\mathbf{P}}^0 \otimes \Phi; \bar{\mathbf{P}}^0 \otimes \Lambda^*) \\ &= \int d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \mathcal{H}(\Phi(s, \mathbf{T}_s, c, \cdot); \mathbf{1}_{\langle c \rangle} \Lambda). \end{aligned}$$

So, the inner entropy density $h^{\text{in}}(\mathbf{P})$ is the conditional relative entropy of its division kernel Φ with respect to Λ^* when the branching tessellation and its cell are selected according to the extended Palm measure $\bar{\mathbf{P}}^0$. In particular, if $h^{\text{in}}(\mathbf{P})$ is finite then the division kernel Φ of \mathbf{P} admits a Radon-Nikodym density with respect to Λ^* .

It is natural to expect that the relative entropy density is affine and lower semi-continuous with compact level sets, at least under some natural caveats. We show this for a topology that is finer than the common weak topology, but is not metrisable. Namely, we define the *topology τ_{loc} of local convergence on \mathcal{P}* as the coarsest topology for which the mapping $\mathbf{P} \mapsto \int f d\mathbf{P}$ is continuous for every bounded local function f . It is then clear that \mathcal{P}_Θ is closed in \mathcal{P} . Recalling the definition (2.31) of the hitting intensity $i_1(\mathbf{P})$, we can then state the following.

Theorem 3.6. *The inner entropy density h^{in} is affine and lower semi-continuous in τ_{loc} . Moreover, for any two constants $0 \leq \beta, \gamma < \infty$ and every $P \in \mathcal{P}_\Theta(\mathbb{T})$, the restricted level set*

$$\mathcal{P}_{\Theta, P, \beta, \gamma} := \left\{ \mathbf{P} \in \mathcal{P}_\Theta : \mathbf{P} \circ \pi_0^{-1} = P, i_1(\mathbf{P}) \leq \beta, h^{\text{in}}(\mathbf{P}) \leq \gamma \right\}$$

is compact and sequentially compact in τ_{loc} .

3.3 Variational principle and existence

Here we change our perspective: Rather than describing a given BRT in terms of its division kernel Φ , we will now suppose that a ‘nice’ division kernel Ψ is given in advance. As we will see, Ψ gives rise to an ‘inner energy’ functional on \mathcal{P}_Θ , and further to an associated ‘inner free energy’, which in turn leads to a variational principle and an existence proof for BRTs with division kernel Ψ . Here are the conditions on Ψ we need.

Definition 3.7. Let us call a division kernel Ψ *moderate* if there exists a measurable density function ψ on the set

$$\{(s, \mathbf{T}_s, c, H) : 0 \leq s \leq 1, \mathbf{T}_s \in \mathbb{B}\mathbb{T}_s, c \in T_s, H \in \langle c \rangle\}$$

satisfying

$$\Psi(s, \mathbf{T}_s, c, dH) = \psi(s, \mathbf{T}_s, c, H) \mathbf{1}_{\langle c \rangle}(H) \Lambda(dH)$$

such that the following holds for all arguments:

- (M1) ψ is *covariant* under translations, in that $\psi(s, \mathbf{T}_s, c, H) = \psi(s, \mathbf{T}_s - x, c - x, H - x)$ for all $x \in \mathbb{R}^d$.
- (M2) ψ has *bounded range*, meaning that there exists a constant $0 \leq r = r_\Psi < \infty$ such that $\psi(s, \mathbf{T}_s, c, H) = \psi(s, \mathbf{T}'_s, c, H)$ whenever $\mathbf{T}_{c+B_r, s} = \mathbf{T}'_{c+B_r, s}$. Here, B_r stands for the closed centred ball with radius r .
- (M3) ψ is *bounded and bounded away from zero*; that is, there exists a constant $\kappa_\Psi < \infty$ such that $|\log \psi| \leq \kappa_\Psi$.
- (M4) Ψ is *approximately STIT for large cells*, which is to say that there exists a constant $\kappa'_\Psi < \infty$ such that

$$\int_{\langle c \rangle} \Lambda(dH) |\psi(s, \mathbf{T}_s, c, H) - 1| \leq \kappa'_\Psi.$$

(In view of the boundedness assumption (M3), this condition involves only the cells c for which $\Lambda(\langle c \rangle)$ is large.)

The assumptions (M1) to (M4) are trivially satisfied if $\Psi = \Lambda^*$ is a STIT kernel as in Example 2.8. More generally, assumptions (M1)–(M3) hold for the cell-driven division kernels in (2.19) whenever the density φ there is uniformly bounded from above and away from zero; (M4) can be achieved by setting $\Psi = \Lambda^*$ for cells with large radius. In Example 2.9, (M1) trivially holds, (M2) holds for each $r > 0$, and (M3) follows from the assumptions on β stated there. Example 2.10 violates the bounded-range property (M2) in the most extreme way conceivable.

A moderate division kernel induces a functional on \mathcal{P}_Θ which, in analogy to the standard Gibbs theory, may be called the (negative) *inner energy in W for Ψ* , and is defined by

$$(3.9) \quad \mathcal{U}_W^{\text{in}}(\mathbf{P}; \Psi) = \int \mathbf{P}(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T}): c \subset W} \log \psi(s, \mathbf{T}_s, c, H);$$

$\mathbf{P} \in \mathcal{P}_\Theta$. (Note that in statistical mechanics the energy always appears negatively in the exponent; so it should not be surprising that $\log \psi$ shows up here. But we suppress the minus sign.) Likewise, there is a term which comes from a normalisation (in our case of the distribution of jump times), and thus may be considered as an analog of the pressure in statistical mechanics. In the present setup, however, this quantity is not only a functional of Ψ , but also of the BRTs $\mathbf{P} \in \mathcal{P}_\Theta$, viz.

$$(3.10) \quad \mathcal{V}_W^{\text{in}}(\mathbf{P}; \Psi) = \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) \sum_{c \in T_s: c \subset W} \int_{\langle c \rangle} \Lambda(dH) (\psi(s, \mathbf{T}_s, c, H) - 1).$$

Theorem 3.8. *For every moderate division kernel Ψ and every $\mathbf{P} \in \mathcal{P}_\Theta$ admitting a covariant division kernel Φ , the following finite limits exist and can be identified:*

$$\begin{aligned} u^{\text{in}}(\mathbf{P}; \Psi) &:= \lim_{n \rightarrow \infty} n^{-d} \mathcal{U}_{[n]}^{\text{in}}(\mathbf{P}_{[n]}; \Psi) = \int \log \psi \, d\bar{\mathbf{P}}^0 \otimes \Phi, \\ v^{\text{in}}(\mathbf{P}; \Psi) &:= \lim_{n \rightarrow \infty} n^{-d} \mathcal{V}_{[n]}^{\text{in}}(\mathbf{P}_{[n]}; \Psi) = \int (\psi - 1) \, d\bar{\mathbf{P}}^0 \otimes \Lambda^*. \end{aligned}$$

In particular, $|u^{\text{in}}(\mathbf{P}; \Psi)| \leq \kappa_\Psi i_1(\mathbf{P})$ and $|v^{\text{in}}(\mathbf{P}; \Psi)| \leq \kappa'_\Psi i_1(\mathbf{P})$.

The energy terms above can be combined with the inner entropy density to define the *inner excess free energy density of \mathbf{P} for Ψ* , namely

$$(3.11) \quad h^{\text{in}}(\mathbf{P}; \Psi) := h^{\text{in}}(\mathbf{P}) - u^{\text{in}}(\mathbf{P}; \Psi) + v^{\text{in}}(\mathbf{P}; \Psi),$$

where the right-hand side is set equal to $+\infty$ if $h^{\text{in}}(\mathbf{P}) = +\infty$. In fact, in the finite case it will turn out that

$$(3.12) \quad h^{\text{in}}(\mathbf{P}; \Psi) = \int d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \mathcal{H}(\Phi(s, \mathbf{T}_s, c, \cdot); \Psi(s, \mathbf{T}_s, c, \cdot)),$$

where Φ is a division kernel for \mathbf{P} . The following variational principle for BRTs is then immediate.

Theorem 3.9. *Let Ψ be any moderate division kernel. A BRT $\mathbf{P} \in \mathcal{P}_\Theta$ then admits Ψ as its division kernel if and only if $h^{\text{in}}(\mathbf{P}; \Psi) = 0$.*

In particular, this can be used to prove the following result.

Theorem 3.10. *For any moderate division kernel Ψ and every $P \in \mathcal{P}_\Theta(\mathbb{T})$, there exists a translation invariant BRT $\mathbf{P} \in \mathcal{P}_\Theta$ with initial distribution P and division kernel Ψ .*

There is a large variety of initial random tessellations P to which this existence theorem applies. The most common examples are the Poisson-Voronoi tessellation, the Poisson-Delaunay tessellation, and the Poisson hyperplane tessellation, which are well-known to satisfy the moment condition (2.38). Further examples are the Delaunay tessellations that are constructed from tempered Gibbsian point processes with tile interaction, as studied in [6]. Unfortunately, we cannot allow a start in a degenerate tessellation with the full space \mathbb{R}^d as its only cell (of any colour), which would be of major interest; cf. the discussion in Example 2.8 on the STIT measure $\Pi^{\Lambda, \infty}$. However and as already indicated above, we can choose the initial distribution P to be the time- δ distribution of $\Pi^{\Lambda, \infty}$ for some small $\delta > 0$. Up to a time shift, this means that there exists a BRT with degenerate start for any moderate division kernel Ψ with an initial cutoff of the form $\psi(s, \cdot, \cdot, \cdot) = 1$ for $0 \leq s < \delta$ and some small δ . In the special case of shape-driven tessellations as in (2.19), the existence of a BRT with degenerate initial tessellation has been proved in [20] under regularity assumptions.

Since $h^{\text{in}}(\cdot; \Psi)$ is affine, the last two theorems imply the following.

Corollary 3.11. *For any moderate division kernel Ψ , the convex set $\mathcal{G}_\Theta(\Psi)$ of all translation invariant BRTs admitting Ψ is a face of \mathcal{P}_Θ . That is, the extremal elements of $\mathcal{G}_\Theta(\Psi)$ are in fact extremal in \mathcal{P}_Θ , and thereby ergodic under translations.*

It is clear that for each ergodic $\mathbf{P} \in \mathcal{G}_\Theta(\Psi)$ its initial distribution $P = \mathbf{P} \circ \pi_0^{-1}$ is also ergodic. The converse holds whenever the correspondence between an initial distribution $P \in \mathcal{P}_\Theta(\mathbb{T})$ and its associated $\mathbf{P} \in \mathcal{G}_\Theta(\Psi)$ is one-to-one. This, however, does not hold in general, as our concluding remark shows.

Remark 3.12. *Uniqueness and phase transition.* It is natural to ask whether or not the convex set $\mathcal{G}(P, \Phi)$ of all BRTs with initial distribution $P \in \mathcal{P}(\mathbb{T})$ and division kernel Φ is a singleton. In general, this is not the case. To provide an example, let $d = 2$ and consider the division kernel Φ defined in Equation (2.20) of Example 2.10. Let $T_{\text{reg}} = \{[1] + i : i \in \mathbb{Z}^2\}$ be the regular tessellation of \mathbb{R}^2 into unit squares and $P \in \mathcal{P}_\Theta(\mathbb{T})$ be given by $P = \int_{[1]} dx \delta_{T_{\text{reg}} - x}$. Further, let \mathbf{P}^{hor} be the STIT tessellation with initial distribution P and driving measure Λ_{hor} as introduced in Example 2.10, and define \mathbf{P}^{vert} analogously. It is then clear that these BRTs live on the spaces \mathbb{T}_{hor} resp. \mathbb{T}_{vert} for all positive times. As a consequence, \mathbf{P}^{hor} and \mathbf{P}^{vert} are two distinct BRTs which both belong to $\mathcal{G}_\Theta(\Phi)$ and have the same initial distribution P .

Although the infinite-range interaction of this example is somewhat artificial, we learn that uniqueness does not hold automatically. Instead, the phenomenon of non-uniqueness, or phase transition, which is a central issue of statistical mechanics, shows up also in the present setting. In analogy to standard results on Gibbs measures (cf.

[9, Section 8.3]), we will show in Proposition 4.17 below that uniqueness does hold for suitable division kernels of bounded range in one spatial dimension. Uniqueness is also known in the non-interacting case (2.19) when the initial tessellation is degenerate and the density φ exhibits some regularity properties [20]. We leave it to the future to find sufficient conditions for uniqueness in higher dimensions as well as examples of bounded-range division kernels exhibiting phase transition. In fact, Figure 2 (right) suggests that a phase transition might already occur for the (moderate) model of Example 2.9.

4 Proofs

4.1 Some properties of local BRTs

Before entering into the proofs of our results we will establish some auxiliary properties of local BRTs. First we will express the local evolution of a BRT in a more explicit form. Throughout this section we let $W \in \mathbb{P}_U$ be an arbitrary window. For any division kernel Φ_W in W we introduce the abbreviation

$$(4.1) \quad \hat{\phi}_W(a, b; \mathbf{T}) = \int_a^b \hat{\phi}_W(s, \mathbf{T}_s) ds,$$

where $0 \leq a < b \leq 1$, $\mathbf{T} \in \mathbb{BT}_W$ and $\hat{\phi}_W(s, \mathbf{T}_s)$ is as in Remark 2.7. For every $T_0 \in \mathbb{T}_W$ we define a measure on \mathbb{BT}_W by

$$(4.2) \quad \begin{aligned} \Pi_W^{\Phi_W}(T_0, \cdot) &= \sum_{n \geq 0} \int \cdots \int_{\{0 \leq s_1 < \cdots < s_n \leq 1\}} ds_1 \cdots ds_n \prod_{i=1}^n \int \hat{\Phi}_W(s_i, \mathbf{T}_{s_{i-1}}, d(c_i, H_i)) \\ &\quad \times \exp \left[-\hat{\phi}_W(0, 1; \mathbf{T}) \right] \mathbb{1}_{\{\pi_0(\mathbf{T})=T_0, \mathcal{D}(\mathbf{T})=\{(s_i, c_i, H_i): 1 \leq i \leq n\}\}} \delta \mathbf{T}, \end{aligned}$$

where $s_0 := 0$ and the last indicator function simply means that \mathbf{T} is the unique branching tessellation which starts from T_0 and is successively defined by the division events (s_i, c_i, H_i) ; recall (2.5).

Lemma 4.1. *Let Φ_W be a division kernel and $\mathbf{P}_W \in \mathcal{P}(\mathbb{BT}_W)$ a BRT in W . Then the following statements are equivalent.*

- (a) \mathbf{P}_W admits the division kernel Φ_W .
- (b) $\Pi_W^{\Phi_W}(T_0, \cdot)$ is the conditional distribution of \mathbf{P}_W given $\pi_0(\mathbf{T}) = T_0$.
- (c) For every non-negative measurable function f on $\overline{\mathbb{BT}_W} \times \mathbb{C}_W \times \langle W \rangle$,

$$(4.3) \quad \begin{aligned} &\int \mathbf{P}_W(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T})} f(s, \mathbf{T}_{s-}, c, H) \\ &= \int_0^1 ds \int \mathbf{P}_{W,s}(d\mathbf{T}_s) \int \hat{\Phi}_W(s, \mathbf{T}_s, d(c, H)) f(s, \mathbf{T}_s, c, H). \end{aligned}$$

- (d) \mathbf{P}_W has the infinitesimal generators $\mathbb{L}_{W,s}^{\Phi_W}$ of (2.15), in that the forward equation

$$(4.4) \quad \int g d\mathbf{P}_{W,t} - \int g d\mathbf{P}_{W,0} = \int_0^t ds \int d\mathbf{P}_{W,s} \mathbb{L}_{W,s}^{\Phi_W} g$$

holds for all $t \in [0, 1]$ and bounded measurable functions g on \mathbb{BT}_W .

Proof. (a) *implies* (b). Recall the recursion steps from Remark 2.7. For $i = 0$, $T_0 = \mathbf{T}_0$ is chosen according to the distribution $\mathbf{P}_{W,0}$. For each $i \geq 1$, conditionally on the first $(i-1)$ division events, the i th division event (s_i, c_i, H_i) for a random tessellation \mathbf{T} with division rule Φ_W is chosen according to the distribution

$$\exp \left[-\hat{\phi}_W(s_{i-1}, s_i; \mathbf{T}) \right] ds_i \hat{\Phi}_W(s_i, \mathbf{T}_{s_{i-1}}, d(c_i, H_i));$$

here we have used that $\mathbf{T}_s = \mathbf{T}_{s_{i-1}}$ for $s \in [s_{i-1}, s_i]$. So, on the event $\{|\mathcal{D}(\mathbf{T})|\}$, the joint distribution of the elements of $\mathcal{D}(\mathbf{T})$ is the product of these conditional measures for $i = 1, \dots, n$, times the probability that s_n is the last division time before 1, which is $\exp \left[-\hat{\phi}_W(s_n, 1; \mathbf{T}) \right]$.

(b) *implies* (c). Fix any initial tessellation T_0 . On the set of all \mathbf{T} with fixed number $|\mathcal{D}(\mathbf{T})| \geq 1$ of division events, the measure $\mathbf{P}_W^{\Phi_W}(T_0, \cdot)$ has a product structure. The elements of $\mathcal{D}(\mathbf{T})$ can be labeled with $i \in \{1, \dots, n\}$ according to their temporal order. For each i we extract the i th term from the product, omit its index i , and separate the terms concerning the division events before and after time $s = s_i$. That is, we write $\mathcal{D}(\mathbf{T}) = \mathcal{D}' \cup \{(s, c, H)\} \cup \mathcal{D}''$ and separate the respective conditional measures. By Fubini's theorem, s can be considered as fixed. For given s , \mathcal{D}' has the same distribution as $\mathcal{D}(\mathbf{T}_s)$, but we still have the condition that the extracted division event (s, c, H) has rank i in $\mathcal{D}(\mathbf{T})$. This condition disappears by summing over i and $n \geq i$. Finally, the division events in \mathcal{D}'' can be integrated out because these do not enter into $f(s, \mathbf{T}_s, c, H)$, and an integration over T_0 gives (c).

(c) *implies* (d). Applying Equation (4.3) to the function

$$f(s, \mathbf{T}, c, H) = \mathbf{1}_{[0,t]}(s) \left[g(\mathcal{O}_{s,c,H}(\mathbf{T}_s)) - g(\mathbf{T}_s) \right],$$

we find that for each $t \in [0, 1]$

$$\begin{aligned} \int_0^t ds \int d\mathbf{P}_{W,s} \mathbb{L}_{W,s}^{\Phi_W} g &= \int \mathbf{P}_W(d\mathbf{T}_W) \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T}_W): s \leq t} [g(\mathbf{T}_{W,s}) - g(\mathbf{T}_{W,s-})] \\ &= \int \mathbf{P}_W(d\mathbf{T}_W) [g(\mathbf{T}_{W,t}) - g(\mathbf{T}_{W,0})] = \int g d\mathbf{P}_{W,t} - \int g d\mathbf{P}_{W,0}. \end{aligned}$$

(d) *implies* (a). In principle, this follows from Reference [7] which, however, makes use of a time-continuity condition on Φ_W . We thus indicate a direct argument. For brevity we omit most indices referring to W . Let $0 \leq s < t \leq 1$ and $\mathbf{P}_{t|\mathbf{T}_s}$ be a regular version of the conditional probability of $\mathbf{P}_{W,t}$ given $\pi_{W,s} = \mathbf{T}_s$. Using (4.4) for a function g of the form $g = \mathbf{1}_A \mathbf{1}_B$ with $A \in \mathcal{B}_{W,s}$ and $B \in \sigma(\mathbf{T} \mapsto (T_{W,u})_{s < u \leq 1})$ and varying A , one readily finds that

$$(4.5) \quad \mathbf{P}_{t|\mathbf{T}_s}(B) - \delta_{\mathbf{T}_s}(B) = \int_s^t du \int d\mathbf{P}_{u|\mathbf{T}_s} \mathbb{L}_u^{\Phi} \mathbf{1}_B$$

for almost all \mathbf{T}_s . We now fix \mathbf{T}_s and think of each \mathbf{T}_u as an element of $\mathbb{B}\mathbb{T}_W$ which is constant on $[u, 1]$. Also, for $\mathbf{T} \in \mathbb{B}\mathbb{T}_W$ we let $\tau(\mathbf{T})$ be the time of the first jump of \mathbf{T} after time s , which is set equal to ∞ when there is no jump during $[s, 1]$. Setting $B = \{\mathbf{T} \in \mathbb{B}\mathbb{T}_W : \tau(\mathbf{T}) > 1\}$, we then find from (4.5) that

$$\mathbf{P}_{t|\mathbf{T}_s}(\tau > t) = 1 - \int_s^t du \hat{\phi}(u, \mathbf{T}_s) \mathbf{P}_{u|\mathbf{T}_s}(\tau > u)$$

and therefore $\mathbf{P}_{t|\mathbf{T}_s}(\tau > t) = \exp[-\hat{\phi}(s, t; \mathbf{T}_s)]$. In other words, τ has the conditional distribution used in Remark 2.7.

Next, let $\Gamma \subset \mathbb{C} \times \mathbb{H}$ be measurable and $\mathcal{O}_\Gamma(T_s) = \{\mathcal{O}_{c,H}(T_s) : c \in T_s, (c, H) \in \Gamma\}$. Consider the set $B = \{\mathbf{T} \in \mathbb{B}\mathbb{T}_W : T_1 \in \mathcal{O}_\Gamma(T_s)\}$ and let $\tau_2(\mathbf{T})$ the time of the second

jump of \mathbf{T} after s (which again is set equal to ∞ if no second jump exists). Then $\mathbf{1}_B(\mathbf{T}_u) = \mathbf{1}_{\{\tau \leq u < \tau_2, T_\tau \in \mathcal{O}_\Gamma(T_s)\}}(\mathbf{T})$ for $u > s$, and (4.5) implies that

$$\begin{aligned} \mathbf{P}_{t|\mathbf{T}_s}(\tau \leq t < \tau_2, T_\tau \in \mathcal{O}_\Gamma(T_s)) &= \int_s^t du \hat{\Phi}(u, \mathbf{T}_s, \Gamma) \mathbf{P}_{u|\mathbf{T}_s}(\tau > u) \\ &\quad - \int_{(s,t] \times \mathcal{O}_\Gamma(T_s)} \mathbf{P}_{t|\mathbf{T}_s}((\tau, \mathbf{T}_\tau) \in d(v, \mathbf{T}_v)) \int_v^t du \hat{\phi}(u, \mathbf{T}_v) \mathbf{P}_{t|\mathbf{T}_v}(\tau > u). \end{aligned}$$

Using the explicit conditional distribution of τ derived above, we thus find that the first term on the right-hand side of the above equation is equal to

$$\int_s^t du \hat{\Phi}(u, \mathbf{T}_s, \Gamma) \exp[-\hat{\phi}(s, u; \mathbf{T}_s)],$$

whereas the second term equals

$$\begin{aligned} \int_{(s,t] \times \mathcal{O}_\Gamma(T_s)} \mathbf{P}_{t|\mathbf{T}_s}((\tau, \mathbf{T}_\tau) \in d(v, \mathbf{T}_v)) \mathbf{P}_{t|\mathbf{T}_v}(v < \tau \leq t) \\ = \mathbf{P}_{t|\mathbf{T}_s}(\tau_2 \leq t, T_\tau \in \mathcal{O}_\Gamma(T_s)). \end{aligned}$$

We thus arrive at the equation

$$\mathbf{P}_{t|\mathbf{T}_s}(\tau \leq t, T_\tau \in \mathcal{O}_\Gamma(T_s)) = \int_s^t du \exp[-\hat{\phi}(s, u; \mathbf{T}_s)] \hat{\Phi}(u, \mathbf{T}_s, \Gamma),$$

which assures that (τ, \mathbf{T}_τ) has the correct conditional distribution of Remark 2.7. \square

Note that the joint integrating measure on the right-hand side of (4.3) can be written in the concise form $\bar{\mathbf{P}}_W \otimes \hat{\Phi}_W$ or, equivalently, $\bar{\mathbf{C}}^{\mathbf{P}_W} \otimes \Phi$, where $\bar{\mathbf{P}}_W = \bar{\mathbf{P}} \circ \pi_W^{-1}$, $\bar{\mathbf{P}}$ is given by (2.11) and $\bar{\mathbf{C}}^{\mathbf{P}_W}$ by (2.36). We will switch between both representations according to convenience.

Corollary 4.2. *Let $\mathbf{P}_W, \mathbf{Q}_W \in \mathcal{P}(\mathbb{BT}_W)$ be two BRTs in W . Suppose \mathbf{Q}_W admits a division kernel Ψ_W , and $\mathbf{P}_W \ll \mathbf{Q}_W$. Then there exists a measurable function $\varphi_W(s, \mathbf{T}_s, c, H) \geq 0$ such that the measure kernel*

$$\Phi_W(s, \mathbf{T}_s, c, dH) := \varphi_W(s, \mathbf{T}_s, c, H) \Psi_W(s, \mathbf{T}_s, c, dH)$$

is a division kernel for \mathbf{P}_W .

Proof. For brevity we introduce the measure kernel

$$\mathbf{D}_W(\mathbf{T}, \cdot) = \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T})} \delta_{(s, \mathbf{T}_{s-}, c, H)}$$

for $\mathbf{T} \in \mathbb{BT}_W$. The integration on the left-hand side of (4.3) is then with respect to the measure $\mathbf{P}_W \mathbf{D}_W$. Since $\mathbf{P}_W \ll \mathbf{Q}_W$ by assumption, it follows that $\mathbf{P}_W \mathbf{D}_W \ll \mathbf{Q}_W \mathbf{D}_W$ with a Radon-Nikodym density f , say. It also follows that $\bar{\mathbf{P}}_W \ll \bar{\mathbf{Q}}_W$ with a density g . Define

$$\varphi_W(s, \mathbf{T}_s, c, H) = f(s, \mathbf{T}_s, c, H) / g(s, \mathbf{T}_s)$$

if the denominator is positive, and zero otherwise. Then we obtain, using equation (4.3) for (\mathbf{Q}_W, Ψ_W) in place of (\mathbf{P}_W, Φ_W) ,

$$\mathbf{P}_W \mathbf{D}_W = f(\mathbf{Q}_W \mathbf{D}_W) = f(\bar{\mathbf{Q}}_W \otimes \hat{\Psi}) = \bar{\mathbf{P}}_W \otimes (\varphi_W \hat{\Psi}_W).$$

In view of Lemma 4.1, this means that \mathbf{P}_W admits the division kernel $\Phi_W := \varphi_W \Psi_W$. \square

Finally we look at the first-moment condition (2.10).

Lemma 4.3. *Let Φ_W be a division kernel for W , and suppose its total mass satisfies the uniform bound (2.14). Then*

$$\int \mathbf{\Pi}_W^{\Phi_W}(T_0, d\mathbf{T}) |T_1| \leq e^\phi |T_0|$$

for all initial tessellations $T_0 \in \mathbb{T}_W$. Moreover, for every $\varepsilon > 0$ one can find a number $\tau < \infty$ such that

$$\int \mathbf{\Pi}_W^{\Phi_W}(T_0, d\mathbf{T}) (|T_1| - \tau |T_0|)_+ \leq \varepsilon |T_0|$$

for all $T_0 \in \mathbb{T}_W$.

Proof. Recall the description of $\mathbf{\Pi}_W^{\Phi_W}(T_0, \cdot)$ in Remark 2.7. The algorithm there implies that, for each i with $s_i \leq 1$, the holding time $s_i - s_{i-1}$ dominates an exponential time with parameter $|T_{s_{i-1}}| \phi$, independently of the previous recursion steps. Hence, the process $|T_s|$ is stochastically dominated by the Furry-Yule process $Z_s \in \mathbb{N}$ with birth rate ϕ , namely the pure birth Markov process which starts in $k = |T_0|$ and jumps from any $j \geq 1$ to $j+1$ with rate $j\phi$. Equivalently, Z_s can be described as the branching process in which each individual, independently of all others, lives for an exponential time with parameter ϕ and then splits into two offspring. In particular, the descendance trees of each of the k ancestors are independent, and it is sufficient to look at the number of descendants at time s in each of these trees. This number is known to have the geometric distribution with mean $e^{\phi s}$. A proof of this can be found e.g. in [16], Examples 6.4, 6.8 or Exercise 6.11.

As for the second assertion, we conclude from the convexity of the function $a \mapsto (a - \tau)_+$ that

$$(|T_1| - \tau |T_0|)_+ \leq \sum_{c \in T_0} (|T_{c,1}| - \tau)_+.$$

Here, $|T_{c,1}|$ is the number of descendants of the initial cell c at time 1, which is stochastically dominated by the geometric random variable Z_1 . As $\mathbb{E}(Z_1 - \tau)_+ \rightarrow 0$ as $\tau \rightarrow \infty$, the result follows immediately. \square

4.2 From local to global division kernels

Let $\mathbf{P} \in \mathcal{P}$ be given and suppose (LAC) holds. Corollary 4.2 then implies that for each $W \in \mathbb{P}$ there exists a cell division kernel Φ_W in W such that \mathbf{P}_W is a BRT for Φ_W . (In fact, Φ_W is absolutely continuous with respect to Λ^* , but we do not need this here.) Therefore, our first step in the proof of Theorem 3.1 will be to construct a global common extension Φ of these kernels Φ_W . Then, in a second step later, we will verify the properties of this Φ .

Proof of Theorem 3.1, first part: Construction of a global division kernel. In the situation above, the measure $\overline{\mathbf{P}}_W \otimes \hat{\Phi}_W$ is concentrated on $\overline{\mathbb{B}\mathbb{T}}_W \times \Delta_W$, where

$$\overline{\mathbb{B}\mathbb{T}}_W = \{(s, \mathbf{T}_{W,s}) : s \in [0, 1], \mathbf{T}_{W,s} \in \mathbb{B}\mathbb{T}_{W,s}\}$$

and $\Delta_W = \{(c, H) : c \in \mathbb{C}_W, H \in \langle c \rangle\}$. Take any event A in $\overline{\mathbb{B}\mathbb{T}}_W \times \Delta_W$ and recall (3.1). Since

$$\mathcal{D}(\mathbf{T}_W) = \{(s, c \cap W, H) : (s, c, H) \in \mathcal{D}(\mathbf{T}), (c, H) \in \tilde{\pi}_W^{-1} \Delta_W\},$$

Lemma 4.1(c) then tells us that

$$(4.6) \quad \overline{\mathbf{P}}_W \otimes \hat{\Phi}_W(A) = \int \mathbf{P}(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T})} \mathbf{1}_{(\tilde{\pi}_W \otimes \tilde{\pi}_W)^{-1} A}(s, \mathbf{T}_{s-}, c, H).$$

As a consequence, the measures $\overline{\mathbf{P}}_W \otimes \hat{\Phi}_W$ with $W \in \mathbb{P}$ are consistent in the sense that

$$(4.7) \quad (\overline{\mathbf{P}}_{W'} \otimes \hat{\Phi}_{W'}) \circ (\overline{\pi}_W \otimes \tilde{\pi}_W)^{-1} = \overline{\mathbf{P}}_W \otimes \hat{\Phi}_W \quad \text{for } W \subset W' \in \mathbb{P}.$$

To define a common extension of these measures we first localise to a fixed window $V \in \mathbb{P}$. For $W \supset V$ we write $\overline{\mathbf{P}}_W \otimes \mathbf{1}_V \hat{\Phi}_W$ for the restriction of $\overline{\mathbf{P}}_W \otimes \hat{\Phi}_W$ to the set $\overline{\mathbb{B}\mathbb{T}}_W \times \tilde{\pi}_V^{-1} \Delta_V$. These measures have a finite total mass that does not depend on W . Indeed, (4.6) and the first-moment condition (2.10) imply that

$$(4.8) \quad \begin{aligned} \|\overline{\mathbf{P}}_W \otimes \mathbf{1}_V \hat{\Phi}_W\| &= \int \mathbf{P}(d\mathbf{T}) |\{(s, c, H) \in \mathcal{D}(\mathbf{T}) : (c, H) \in \tilde{\pi}_V^{-1} \Delta_V\}| \\ &\leq \int \mathbf{P}(d\mathbf{T}) |T_{V,1}| < \infty. \end{aligned}$$

Since all spaces under consideration are Borel spaces, we can thus apply an abstract version of the Kolmogorov extension theorem [11, Corollary 6.15] to obtain a finite measure on $\overline{\mathbb{B}\mathbb{T}} \times \tilde{\pi}_V^{-1} \Delta_V$, to be denoted by $\overline{\mathbf{P}} \otimes \mathbf{1}_V \hat{\Phi}$, which satisfies

$$(\overline{\mathbf{P}} \otimes \mathbf{1}_V \hat{\Phi}) \circ (\overline{\pi}_W \otimes \tilde{\pi}_W)^{-1} = \overline{\mathbf{P}}_W \otimes \mathbf{1}_V \hat{\Phi}_W$$

for all $W \in \mathbb{P}$ with $W \supset V$. Since V is arbitrary and

$$\bigcup_{V \in \mathbb{P}} \tilde{\pi}_V^{-1} \Delta_V = \Delta := \{(c, H) \in \mathbb{C} \times \mathbb{H} : H \in \langle c \rangle\},$$

the measures $\overline{\mathbf{P}} \otimes \mathbf{1}_V \hat{\Phi}$ can be glued together to a locally finite measure $\overline{\mathbf{P}} \otimes \hat{\Phi}$ on $\overline{\mathbb{B}\mathbb{T}} \times \Delta$ satisfying

$$(4.9) \quad (\overline{\mathbf{P}} \otimes \hat{\Phi}) \circ (\overline{\pi}_W \otimes \tilde{\pi}_W)^{-1} = \overline{\mathbf{P}}_W \otimes \hat{\Phi}_W \quad \text{for all } W \in \mathbb{P}.$$

As we have indicated by the notation, disintegration shows that this measure is indeed the product of $\overline{\mathbf{P}}$ with a locally finite measure kernel $\hat{\Phi}$.

We must show that this $\hat{\Phi}$ is indeed a (cumulative) division kernel. By construction, each $\hat{\Phi}(s, \mathbf{T}_s, \cdot)$ is supported on Δ , which means that each $\Phi(s, \mathbf{T}_s, c, \cdot)$ is supported on $\langle c \rangle$. In fact, considering the set $\Delta(T_s) = \{(c, H) : c \in T_s, H \in \langle c \rangle\}$ and its complement $\neg \Delta(T_s)$, we can write

$$\begin{aligned} &\int d\overline{\mathbf{P}}(s, \mathbf{T}_s) \hat{\Phi}(s, \mathbf{T}_s, \neg \Delta(T_s)) \\ &= \int d(\overline{\mathbf{P}} \otimes \hat{\Phi})(s, \mathbf{T}_s, c, H) \lim_{W \uparrow \mathbb{R}^d} (1 - \mathbf{1}_{\tilde{\pi}_W^{-1} \Delta(T_{W,s})}(c, H)), \end{aligned}$$

and the last term vanishes by (4.9) and Fatou's lemma. So we can conclude that, for $\overline{\mathbf{P}}$ -almost all (s, \mathbf{T}_s) , $\hat{\Phi}(s, \mathbf{T}_s, \cdot)$ is indeed supported on $\Delta(T_s)$, as required. \square

It remains to prove that the kernel Φ satisfying (4.9) satisfies the properties (a) to (d) of Theorem 3.1. Observe first that statement (d) follows directly from (4.9). As for (a), a combination of (4.6) and (4.9) shows that the asserted identity holds when there is some $W \in \mathbb{P}$ and a measurable function f_W with

$$f(s, \mathbf{T}, c, H) = \mathbf{1}_{\langle c \cap W \rangle}(H) f_W(s, \mathbf{T}_W, c \cap W, H).$$

The general identity thus follows by a monotone class argument. Furthermore, the proof of (b) can be literally taken from the proof of the implication '(c) implies (d)' of Lemma 4.1.

The proof of the Gibbs property (c) requires some more work. Let us first introduce a modification of the outer projection for a given window $W \in \mathbb{P}$, which refers to a

larger window $W' \in \mathbb{P}$ rather than the full space \mathbb{R}^d . Namely, for $W \subset W' \in \mathbb{P}$ and any $\mathbf{T}_{W'} \in \mathbb{B}\mathbb{T}_{W'}$ let

$$(4.10) \quad \mathbf{T}_W^{W', \text{out}} = (\{c \in T_{W', s} : c \notin \text{int}(W)\})_{0 \leq s \leq 1}$$

be the evolution of the cells hitting $W' \setminus \text{int}(W)$. In particular, if $W' = W$ then $\mathbf{T}_W^{W, \text{out}} = \mathbf{T}_W^\partial$. We also set $\mathbf{T}_W^{W', 0, \text{out}} = (T_{W, 0}^{\text{in}}, \mathbf{T}_W^{W', \text{out}})$ and let $\mathcal{B}_W^{W', 0, \text{out}}$ denote the σ -field on $\mathbb{B}\mathbb{T}_{W'}$ generated by the mapping $\mathbf{T}_{W'} \mapsto \mathbf{T}_W^{W', 0, \text{out}}$.

Proof of the Gibbs property, Theorem 3.1(c). Fix a window $W \in \mathbb{P}$ and let $\mathbf{P}_W^{\text{in}}(\cdot | \mathbf{T}_W^{0, \text{out}})$ be a regular version of the conditional distribution of π_W^{in} under the condition $\pi_W^{0, \text{out}} = \mathbf{T}_W^{0, \text{out}}$. We need to show that this probability kernel coincides with $\mathbf{G}_W^\Phi(\cdot | \mathbf{T}_W^{0, \text{out}})$ almost surely. (In particular, this will imply that the latter is almost surely well-defined.) Pick any two non-negative measurable functions $g(\mathbf{T}_W^{0, \text{out}})$ and $h(s, \mathbf{T}_{W, s}^{\text{in}}, c, H)$ of the indicated arguments. We suppose g is local, in that $g(\mathbf{T}_W^{0, \text{out}}) = g(\mathbf{T}_W^{W', 0, \text{out}})$ for some $W' \in \mathbb{P}$ containing W . Consider the integral

$$(4.11) \quad \int \mathbf{P}(d\mathbf{T}) g(\mathbf{T}_W^{0, \text{out}}) \int \mathbf{P}_W^{\text{in}}(d\mathbf{S} | \mathbf{T}_W^{0, \text{out}}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{S})} h(s, \mathbf{S}_{s-}, c, H).$$

By the definition of conditional distribution, this is equal to

$$\int \mathbf{P}(d\mathbf{T}) g(\mathbf{T}_W^{0, \text{out}}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T}_W^{\text{in}})} h(s, \mathbf{T}_{W, s-}^{\text{in}}, c, H).$$

In view of the locality assumption on g , the integrand actually only depends on $\mathbf{T}_{W'}$, which is a pure jump process and therefore strongly Markov. Writing (s_i, c_i, H_i) for the i th division event of $\mathbf{T}_{W'}$ in temporal order, we can rewrite the last expression in the form

$$(4.12) \quad \sum_{i \geq 1} \int \mathbf{P}_{W'}(d\mathbf{T}_{W'}) g(\mathbf{T}_W^{W', 0, \text{out}}) \mathbf{1}_{\{s_i < 1, c_i \subset W\}}(\mathbf{T}_{W'}) h(s_i, \mathbf{T}_{W, s_i-}^{\text{in}}, c_i, H_i).$$

Now, both h and the indicator function in the integrand are measurable with respect to the σ -field \mathcal{B}_{W', s_i} of all events $A \in \mathcal{B}_{W'}$ with $A \cap \{s_i \leq t\} \in \mathcal{B}_{W', t}$ for all t . By the strong Markov property, we can therefore replace the function $g(\mathbf{T}_W^{W', 0, \text{out}})$ by its conditional expectation $g(s_i, \mathbf{T}_{W', s_i})$ relative to \mathcal{B}_{W', s_i} . Furthermore, the process $\mathbf{T}_W^{W', 0, \text{out}}$ is itself a Markov jump process. (In fact, it can be considered as the BRT in W' for the division kernel which equals $\Phi_{W'}(s, \mathbf{T}_{W'}, c, \cdot)$ if $c \not\subset W$ and is identically zero otherwise.) This means that

$$g(s_i, \mathbf{T}_{W', s_i}) = g(s_i, \mathbf{T}_{W, s_i}^{W', 0, \text{out}}) = g(s_i, \mathbf{T}_{W', s_i-})$$

when $c_i \subset W$. Altogether, we find that the expression (4.12) is equal to

$$\int \mathbf{P}(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T}) : c \subset W} g(s, \mathbf{T}_{W', s-}) h(s, \mathbf{T}_{W, s-}^{\text{in}}, c, H).$$

By statement (a), this in turn coincides with

$$\int_0^1 ds \int \mathbf{P}_s(d\mathbf{T}_s) \sum_{c \in T_s : c \subset W} \int \Phi(s, \mathbf{T}_s, c, dH) g(s, \mathbf{T}_{W', s}) h(s, \mathbf{T}_{W, s}^{\text{in}}, c, H),$$

which by the Markov property is equal to

$$\int_0^1 ds \int \mathbf{P}(d\mathbf{T}) g(\mathbf{T}_W^{0,\text{out}}) \sum_{c \in T_{W,s}^{\text{in}}} \int \Phi(s, \mathbf{T}_s, c, dH) h(s, \mathbf{T}_{W,s}^{\text{in}}, c, H).$$

Taking conditional expectation with respect to $\mathbf{T}_W^{0,\text{out}}$ and using the conditional division kernel $\Phi_W^{\text{in}}(\cdot | \mathbf{T}_W^{\text{out}})$ from Definition 2.11, we can rewrite this as

$$\begin{aligned} \int \mathbf{P}(d\mathbf{T}) g(\mathbf{T}_W^{0,\text{out}}) \int_0^1 ds \int \mathbf{P}_{W,s}^{\text{in}}(d\mathbf{S}_s | \mathbf{T}_W^{0,\text{out}}) \\ \times \int \hat{\Phi}_W^{\text{in}}(s, \mathbf{S}_s, c, dH | \mathbf{T}_W^{\text{out}}) h(s, \mathbf{S}_s, c, H). \end{aligned}$$

Since the underlying spaces are Borel, a comparison of (4.11) with the preceding expression shows that, for almost all $\mathbf{T}_W^{0,\text{out}}$,

$$\bar{\mathbf{P}}_W^{\text{in}}(\cdot | \mathbf{T}_W^{0,\text{out}}) \otimes \hat{\Phi}_W^{\text{in}}(\cdot | \mathbf{T}_W^{\text{out}}) = \int \mathbf{P}_W^{\text{in}}(d\mathbf{S} | \mathbf{T}_W^{0,\text{out}}) \sum_{(s,c,H) \in \mathcal{D}(\mathbf{S})} \delta_{(s, \mathbf{S}_{s-}^{\text{in}}, c, H)},$$

which corresponds to (4.3). Lemma 4.1 therefore implies that $\mathbf{P}_W^{\text{in}}(\cdot | \mathbf{T}_W^{0,\text{out}})$ coincides with $\mathbf{G}_W^{\Phi}(\cdot | \mathbf{T}_W^{0,\text{out}})$. This completes the proof of the Gibbs property (c). \square

Here are two supplements to the preceding proofs. The first deals with the consistency properties (4.7) resp. (4.9), and the second with a localised version of the Gibbs property (c).

Remark 4.4. *Consistency of kernel densities.* Consider two windows $W, W' \in \mathbb{P}$ with $W \subset W'$ and a BRT $\mathbf{P} \in \mathcal{P}$ satisfying $\mathbf{P}_{W'} \ll \mathbf{P}_{W',0} \Pi_{W'}^{\Lambda}$. Let φ_W and $\varphi_{W'}$ be the Λ -densities of the division kernels Φ_W and $\Phi_{W'}$ of \mathbf{P}_W and $\mathbf{P}_{W'}$, which exist by Corollary 4.2. The consistency equation (4.7) then means that

$$\varphi_W(s, \mathbf{T}_{W,s}, c, H) = \int \mathbf{P}_{W',s | \mathbf{T}_{W,s}}(d\mathbf{T}_{W',s}) \varphi_{W'}(s, \mathbf{T}_{W',s}, \pi_W^{-1}(c, T_{W',s}), H)$$

for $\bar{\mathbf{P}}_W \otimes \Lambda_W^*$ almost all arguments. Here, $\mathbf{P}_{W',s | \mathbf{T}_{W,s}}$ stands for a regular version of the conditional distribution of $\mathbf{T}_{W',s}$ under \mathbf{P}_s given $\mathbf{T}_{W,s}$, and $c' = \pi_W^{-1}(c, T_{W',s})$ is the unique element of $T_{W',s}$ with $c' \cap W = c$. An analogous statement holds for $W' = \mathbb{R}^d$ when \mathbf{P} admits a global division kernel with a Λ -density.

Remark 4.5. *Localising the outer conditions for conditional BRTs.* Fix two windows $W, W' \in \mathbb{P}$ with $W \subset W'$ and let $\mathbf{P}_{W'}$ be a BRT in W' for a division kernel $\Phi_{W'}$. Furthermore, replace \mathbb{R}^d by W' in Definition 2.11 and use the conditional division kernel

$$\Phi_W^{\text{in}}(s, \mathbf{S}_s, c, \cdot | \mathbf{T}_{W'}) := \Phi_{W'}(s, \mathbf{S}_s \cup \mathbf{T}_{W,s}^{W',\text{out}}, c, \cdot),$$

to obtain a conditional BRT $\mathbf{G}_W^{\Phi_{W'}}(\cdot | \mathbf{T}_W^{W',0,\text{out}})$ in W ; here we use the notation introduced in and after (4.10). The arguments in the proof of Theorem 3.1(c) then show that the kernel $\mathbf{G}_W^{\Phi_{W'}}$ is a regular version of the conditional distribution of π_W^{in} for $\mathbf{P}_{W'}$ under the condition $\mathcal{B}_W^{W',0,\text{out}}$.

Next we show that, for translation invariant \mathbf{P} , Φ can be chosen to be covariant under translations.

Proof of Corollary 3.3. We only consider the covariance of the global division kernel. The second statement of the corollary is then obvious from Theorem 3.1(d). Let $\mathbf{P} \in \mathcal{P}_\Theta$ be as stated and Φ a global division kernel for \mathbf{P} , which exists by Theorem 3.1. For $x \in \mathbb{R}^d$ let ϑ_x be the spatial translation by $-x$, which acts on $\overline{\mathbb{BT}}$ via $\vartheta_x : (s, \mathbf{T}_s) \mapsto (s, \mathbf{T}_s - x)$ and, by hypothesis, leaves $\overline{\mathbf{P}}$ invariant. As before, we use the same symbol for the translation $\vartheta_x : (c, H) \mapsto (c - x, H - x)$ acting on $\mathbb{C} \times \mathbb{H}$. For $x \in \mathbb{R}^d$ let $\hat{\Phi}^x(s, \mathbf{T}_s, \cdot) := \hat{\Phi}(s, \mathbf{T}_s + x, \vartheta_x^{-1} \cdot)$. We first claim that

$$(4.13) \quad \overline{\mathbf{P}} \otimes \hat{\Phi} = \overline{\mathbf{P}} \otimes \hat{\Phi}^x$$

for all x . Indeed, let $A \in \mathcal{B}(\overline{\mathbb{BT}})$ and $B \in \mathcal{B}(\mathbb{C} \times \mathbb{H})$ be arbitrarily given. Then we can write, using the ϑ_x -invariance of \mathbf{P} in the first step,

$$\begin{aligned} (\overline{\mathbf{P}} \otimes \hat{\Phi}^x)(A \times B) &= (\overline{\mathbf{P}} \otimes \hat{\Phi})(\vartheta_x^{-1} A \times \vartheta_x^{-1} B) \\ &= \int \mathbf{P}(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T}) : (c-x, H-x) \in B} \mathbf{1}_A(s, \mathbf{T}_{s-}) \\ &= \int \mathbf{P}(d\mathbf{T}) \sum_{(s, c, H) \in \mathcal{D}(\mathbf{T}) : (c, H) \in B} \mathbf{1}_A(s, \mathbf{T}_{s-}) \\ &= (\overline{\mathbf{P}} \otimes \hat{\Phi})(A \times B), \end{aligned}$$

proving (4.13). The second and the fourth step come from Theorem 3.1(a), and in the third step we observed that $\mathcal{D}(\mathbf{T})$ consists of the shifted elements of $\mathcal{D}(\mathbf{T} - x)$ and then used again the translation invariance of \mathbf{P} . Equation (4.13) shows that $\hat{\Phi}^x(\cdot, \cdot, B) = \hat{\Phi}(\cdot, \cdot, B)$ $\overline{\mathbf{P}}$ -almost surely for each B .

To obtain an everywhere covariant version of Φ , we pick a countable generator \mathcal{G} of $\mathcal{B}(\mathbb{C} \times \mathbb{H})$ which is stable under intersections. We also let Γ be the set of all $(s, \mathbf{T}_s) \in \overline{\mathbb{BT}}$ which are such that $\hat{\Phi}(s, \mathbf{T}_s, B) = \hat{\Phi}^y(s, \mathbf{T}_s, B)$ for all $B \in \mathcal{G}$ (and thus all $B \in \mathcal{B}(\mathbb{C} \times \mathbb{H})$) and the countably many lattice elements $y \in \mathbb{Z}^d$. First, (4.13) implies that $\overline{\mathbf{P}}(\Gamma) = 1$. As an intermediate step, we then define

$$\tilde{\Phi}(s, \mathbf{T}_s, \cdot) = \begin{cases} \hat{\Phi}(s, \mathbf{T}_s, \cdot) & \text{if } (s, \mathbf{T}_s) \in \Gamma, \\ \hat{\Lambda}^*(s, \mathbf{T}_s, \cdot) & \text{otherwise,} \end{cases}$$

where $\hat{\Lambda}^*$ is the cumulative STIT kernel of Example 2.8. It is then clear that $\tilde{\Phi}$ is a version of $\hat{\Phi}$ which satisfies $\tilde{\Phi}^y = \tilde{\Phi}$ for all $y \in \mathbb{Z}^d$.

To achieve the covariance under the full translation group, we finally define

$$\bar{\Phi} = \int_{[1]} \tilde{\Phi}^x dx,$$

where $[1]$ is the centred unit cube in \mathbb{R}^d . Then for each $y \in \mathbb{R}^d$ we have

$$\bar{\Phi}^y = \int_{[1]+y} \tilde{\Phi}^x dx = \bar{\Phi}$$

because $[1]+y$ can be decomposed into finitely many pieces which are lattice translations of corresponding pieces of $[1]$. On the other hand, since

$$\overline{\mathbf{P}} \otimes \bar{\Phi} = \int_{[1]} \overline{\mathbf{P}} \otimes \tilde{\Phi}^x dx = \int_{[1]} \overline{\mathbf{P}} \otimes \hat{\Phi}^x dx = \overline{\mathbf{P}} \otimes \hat{\Phi}$$

by (4.13), $\bar{\Phi}$ is also a version of $\hat{\Phi}$. □

4.3 On the inner entropy density

We first recall some standard properties of relative entropy. A basic fact is the variational formula, which states that

$$(4.14) \quad \mathcal{H}(\mu, \nu) = \sup_g \left[\int g \, d\mu - \log \int e^g \, d\nu \right],$$

for any two probability measures μ, ν on a common measurable space. Here, the supremum extends over all bounded measurable functions on this space; see [26, Theorem 4.1]. On the one hand, the variational formula implies the useful estimate

$$(4.15) \quad \int g \, d\mu \leq \mathcal{H}(\mu, \nu) + \log \int e^g \, d\nu$$

for any non-negative measurable g . On the other hand, using Jensen's inequality it follows immediately that $\mathcal{H}(\mu, \nu)$ is jointly measure convex in both arguments simultaneously. Also, it is jointly lower semicontinuous in (μ, ν) in the topology generated by the integrals of bounded measurable functions. Finally, if μ and ν are restricted to a sub- σ -field \mathcal{A} then relative entropy is increasing in \mathcal{A} . Alternative proofs of these facts can be found in [9], Section 15.1, for example. Since $\mathcal{H}(a\mu; b\nu) = a\mathcal{H}(\mu, \nu) + b\varrho(a/b)$ for $a, b > 0$ and normalised μ, ν (recall (3.5)), the last facts extend directly to the case of finite measures, except that the convexity then holds in the first argument only.

Now, turning to the proof of Theorem 3.5 we proceed with a series of lemmas. Let $\mathbf{P} \in \mathcal{P}_\Theta$ be arbitrarily given and $P = \mathbf{P} \circ \pi_0^{-1}$ its initial distribution. We can clearly assume that

$$(4.16) \quad \liminf_{n \rightarrow \infty} n^{-d} \mathcal{H}_{[n]}^{\text{in}}(\mathbf{P}) < \infty$$

because otherwise there is nothing to show.

Lemma 4.6. *Condition (4.16) implies condition (LAC).*

Proof. We fix a window $W \in \mathbb{P}$. For any n with $W \subset [n]$ let

$$(4.17) \quad A_n = \{ \mathbf{T} \in \mathbb{BT} : c \cap W = \emptyset \text{ for all } c \in T_0 \text{ with } c \cap \partial[n] \neq \emptyset \}.$$

Obviously, $A_n \in \mathcal{B}_{[n]}^\partial$. Also, since for each $T \in \mathbb{T}$ the union of all cells hitting W is contained in some $[n]$, we have $A_n \uparrow \mathbb{BT}$ as $n \rightarrow \infty$. Furthermore, if $\mathbf{T} \in A_n$ then the inner window $\text{in}_{[n]}(s, \mathbf{T}_{[n]}^\partial)$ (defined in (2.29)) contains W for all s , so that

$$\mathbf{G}_{[n]}^\Lambda(B | \mathbf{T}_{[n]}^{0, \partial}) = \Pi_W^\Lambda(T_{W, 0}, B)$$

for all $B \in \mathcal{B}_W$. Suppose now that $\mathcal{H}_{[n]}^{\text{in}}(\mathbf{P}) < \infty$ and recall the notation of Definition 3.4. Writing $\mathbf{P}_{[n]}^{\text{in}}(\cdot | \mathbf{T}_{[n]}^{0, \partial})$ for a regular conditional distribution of $\pi_{[n]}^{\text{in}}$ under the condition $\pi_{[n]}^{0, \partial} = \mathbf{T}_{[n]}^{0, \partial}$, we then have

$$\mathbf{P}_{[n]}^{\text{in}}(\cdot | \mathbf{T}_{[n]}^{0, \partial}) \ll \mathbf{G}_{[n]}^\Lambda(\cdot | \mathbf{T}_{[n]}^{0, \partial}) \quad \text{for } \mathbf{P}\text{-almost all } \mathbf{T}.$$

Therefore, if $B \in \mathcal{B}_W$ is such that $P\Pi^\Lambda(B) = 0$ then $\Pi_W^\Lambda(T_{W, 0}, B) = 0$ for almost all \mathbf{T} , and thus $\mathbf{P}_{[n]}^{\text{in}}(B | \mathbf{T}_{[n]}^{0, \partial}) = 0$ for almost all $\mathbf{T} \in A_n$. Hence

$$\mathbf{P}(B \cap A_n) = \int_{A_n} \mathbf{P}(d\mathbf{T}) \mathbf{P}_{[n]}^{\text{in}}(B | \mathbf{T}_{[n]}^{0, \partial}) = 0.$$

Letting $n \rightarrow \infty$ through the integers n with $\mathcal{H}_{[n]}^{\text{in}}(\mathbf{P}) < \infty$, we thus obtain that $\mathbf{P}(B) = 0$. So, we have shown that $\mathbf{P} \ll P\Pi^\Lambda$ on \mathcal{B}_W , and the proof is complete. \square

Combining the preceding lemma with Theorem 3.1, we can conclude that \mathbf{P} admits a global division kernel Φ . Hence, for each window $W \in \mathbb{P}$, the conditional distribution of π_W^{in} given $\mathcal{B}_W^{0,\partial}$ under \mathbf{P} resp. $P\Pi^\Lambda$ are equal to the localised conditional BRTs $\mathbf{G}_W^{\Phi_W}(\cdot | \mathbf{T}_W^{0,\partial})$ resp. $\mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\partial})$ introduced in Remark 4.5 resp. Example 2.12. It follows that

$$(4.18) \quad \mathcal{H}_W^{\text{in}}(\mathbf{P}) = \int \mathbf{P}(\mathrm{d}\mathbf{T}) \mathcal{H}(\mathbf{G}_W^{\Phi_W}(\cdot | \mathbf{T}_W^{0,\partial}); \mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\partial})).$$

This expression can be specified as follows.

Lemma 4.7. *Under (4.16), we have for each $W \in \mathbb{P}$,*

$$(4.19) \quad \mathcal{H}_W^{\text{in}}(\mathbf{P}) = \int \mathrm{d}\bar{\mathbf{P}}_W(s, \mathbf{T}_W) \sum_{c \in T_{W,s}^{\text{in}}} \mathcal{H}(\Phi_W(s, \mathbf{T}_{W,s}, c, \cdot); \mathbf{1}_{\langle c \rangle} \Lambda).$$

Proof. By Lemma 4.6 and Corollary 4.2, Φ_W is absolutely continuous with respect to Λ ; we write φ_W for the associated Radon–Nikodym density. Using the equivalence of Lemma 4.1(a) and (b) separately for the intervals between the ‘immigration times’ (2.30), we obtain the following identity for the Radon–Nikodym density of $\mathbf{G}_W^{\Phi_W}(\cdot | \mathbf{T}_W^{0,\partial})$ relative to $\mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\partial})$:

$$(4.20) \quad \log \frac{\mathrm{d}\mathbf{G}_W^{\Phi_W}(\cdot | \mathbf{T}_W^{0,\partial})}{\mathrm{d}\mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\partial})}(\mathbf{T}_W^{\text{in}}) = \left[\hat{\lambda}_W^{\text{in}}(0, 1; \mathbf{T}_W^{\text{in}}) - \hat{\phi}_W^{\text{in}}(0, 1; \mathbf{T}_W) \right] + \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T}_W^{\text{in}})} \log \varphi_W(s, \mathbf{T}_{W,s-}, c, H),$$

where $\mathbf{T}_W \in \mathbb{B}\mathbf{T}_W$, $\mathcal{D}(\mathbf{T}_W^{\text{in}})$ is the associated set of division events with $c \subset W$, and similarly to (4.1),

$$\hat{\phi}_W^{\text{in}}(0, 1; \mathbf{T}_W) = \int_0^1 \mathrm{d}s \sum_{c \in T_{W,s}^{\text{in}}} \Phi_W(s, \mathbf{T}_{W,s}, c, \langle c \rangle)$$

and $\hat{\lambda}_W^{\text{in}}(0, 1; \mathbf{T}_W^{\text{in}}) = \int_0^1 \mathrm{d}s \sum_{c \in T_{W,s}^{\text{in}}} \Lambda(\langle c \rangle)$.

Now, $\mathcal{H}_W^{\text{in}}(\mathbf{P})$ is simply the integral of (4.20) over $\mathbf{P}_W^{0,\partial}(\mathrm{d}\mathbf{T}_W^{0,\partial}) \mathbf{G}_W^{\Phi_W}(\mathrm{d}\mathbf{T}_W^{\text{in}} | \mathbf{T}_W^{0,\partial}) = \mathbf{P}_W(\mathrm{d}\mathbf{T}_W)$. By the equation in Lemma 4.1(c), integration of the last term in (4.20) gives the contribution

$$\int \mathrm{d}\bar{\mathbf{P}}_W(s, \mathbf{T}_{W,s}) \sum_{c \in T_{W,s}^{\text{in}}} \int_{\langle c \rangle} \Phi_W(s, \mathbf{T}_{W,s}, c, \mathrm{d}H) \log \varphi_W(s, \mathbf{T}_{W,s}, c, H).$$

In terms of the measure $\bar{\mathbf{C}}^{\mathbf{P}_W, \text{in}}$ which is defined by restricting the sum in (2.36) to the cells $c \in T_{W,s}^{\text{in}}$, this can be rewritten in the concise form $\int \mathrm{d}\bar{\mathbf{C}}^{\mathbf{P}_W, \text{in}} \otimes \Lambda^* \varphi_W \log \varphi_W$. Likewise, we have

$$\int \mathbf{P}_W(\mathrm{d}\mathbf{T}_W) \left[\hat{\lambda}_W^{\text{in}}(0, 1; \mathbf{T}_W^{\text{in}}) - \hat{\phi}_W^{\text{in}}(0, 1; \mathbf{T}_W) \right] = \int \mathrm{d}\bar{\mathbf{C}}^{\mathbf{P}_W, \text{in}} \otimes \Lambda^* [1 - \varphi_W].$$

Consequently, the \mathbf{P}_W -integral of (4.20) is equal to $\int \mathrm{d}\bar{\mathbf{C}}^{\mathbf{P}_W, \text{in}} \otimes \Lambda^* \varrho(\varphi_W)$, and (4.19) follows by recalling (3.4). \square

The final step in the proof of Theorem 3.5 is the following. Let $h^{\text{in}}(\mathbf{P})$ be defined by (3.8). For brevity we write $h_n^{\text{in}}(\mathbf{P})^{-d} \mathcal{H}_{[n]}^{\text{in}}(\mathbf{P})$.

Lemma 4.8. Under (4.16),

$$\lim_{n \rightarrow \infty} h_n^{\text{in}}(\mathbf{P}) = \sup_{n \geq 1} h_n^{\text{in}}(\mathbf{P}) = h^{\text{in}}(\mathbf{P}).$$

Proof. We claim first that $\liminf_{n \rightarrow \infty} h_n^{\text{in}}(\mathbf{P}) \geq h^{\text{in}}(\mathbf{P})$. We pick any $0 < \eta < 1$ and $\ell < \infty$ and restrict the sum in (4.19) for $W = [n]$ to the cells of T_s with midpoint in $[\eta n]$ and radius at most ℓ . More precisely, we let $L_n \subset \mathbb{R}^d$ be such that the set $\{[\eta] + x : x \in L_n\}$ is a tessellation of the cube $[n\eta]$. (Note that $|L_n|^d$.) Also, we take any $0 < \varepsilon < 1 - \eta$ and let n be so large that $\eta + \ell < \varepsilon n$. Then we can write

$$h_n^{\text{in}}(\mathbf{P}) \geq n^{-d} \sum_{x \in L_n} h_{n,\eta,\ell}(x)$$

with

$$h_{n,\eta,\ell}(x) := \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) \sum_{c \in T_s \cap \Gamma_{\eta,\ell}(x)} \mathcal{H}(\Phi_{[n]}(s, \mathbf{T}_{[n],s}, c, \cdot); \mathbb{1}_{\langle c \rangle} \Lambda),$$

where $\Gamma_{\eta,\ell}(x)$ is the set of all cells c satisfying $m(c) \in [\eta] + x$ and $r(c) \leq \ell$. Now, whether or not a cell $c \in T_{[n],s}$ belongs to $\Gamma_{\eta,\ell}(x)$ can be decided by looking at the restriction $T_{[\varepsilon n]+x,s}$. So, using Remark 4.4 and Jensen's inequality together with the translation invariance of \mathbf{P} and the covariance equation (3.3) we find for each $x \in L_n$,

$$\begin{aligned} h_{n,\eta,\ell}(x) &= \int d\bar{\mathbf{P}}_{[\varepsilon n]+x}(s, \mathbf{T}_{[\varepsilon n]+x,s}) \sum_{c \in T_s \cap \Gamma_{\eta,\ell}(x)} \int \mathbf{P}_{[n],s | \mathbf{T}_{[\varepsilon n]+x,s}}(d\mathbf{T}_{[n],s}) \\ &\quad \times \int_{\langle c \rangle} \Lambda(dH) \varrho(\varphi_{[n]}(s, \mathbf{T}_{[n],s}, c, H)) \\ &\geq \int d\bar{\mathbf{P}}_{[\varepsilon n]}(s, \mathbf{T}_{[\varepsilon n],s}) \sum_{c \in T_s \cap \Gamma_{\eta,\ell}} \int_{\langle c \rangle} \Lambda(dH) \varrho(\varphi_{[\varepsilon n]}(s, \mathbf{T}_{[\varepsilon n],s}, c, H)) \\ &= \mathcal{H}(\mathbb{1}_{\Gamma_{\eta,\ell}} \bar{\mathbf{C}}^{\mathbf{P}} \otimes \Phi|_{\mathcal{B}_{[\varepsilon n]}}; \mathbb{1}_{\Gamma_{\eta,\ell}} \bar{\mathbf{C}}^{\mathbf{P}} \otimes \Lambda^*|_{\mathcal{B}_{[\varepsilon n]}}). \end{aligned}$$

In the last expression, $\mathcal{B}_{[\varepsilon n]}$ is identified with the σ -field that is generated by the projection $\bar{\pi}_{[\varepsilon n]} \otimes \text{id}$, and $\Gamma_{\eta,\ell} := \Gamma_{\eta,\ell}(0)$ is viewed as a set in the product space $\mathbb{BT} \times \mathbb{C} \times \mathbb{H}$. In the limit as $n \rightarrow \infty$, Perez' continuity theorem for relative entropies (cf. [9, Proposition 15.6]) implies that the last relative entropy converges to

$$\mathcal{H}(\mathbb{1}_{\Gamma_{\eta,\ell}} \bar{\mathbf{C}}^{\mathbf{P}} \otimes \Phi; \mathbb{1}_{\Gamma_{\eta,\ell}} \bar{\mathbf{C}}^{\mathbf{P}} \otimes \Lambda^*) = \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) \sum_{c \in T_s \cap \Gamma_{\eta,\ell}} \mathcal{H}(\Phi(s, \mathbf{T}_s, c, \cdot); \mathbb{1}_{\langle c \rangle} \Lambda).$$

By the time-integrated version of the Palm formula (2.33) and the shift covariance of Φ and Λ , the last integral is equal to

$$h_{\eta,\ell} := \eta^d \int_{\{r(c) \leq \ell\}} d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \mathcal{H}(\Phi(s, \mathbf{T}_s, c, \cdot); \mathbb{1}_{\langle c \rangle} \Lambda).$$

Altogether, we find that $\liminf_{n \rightarrow \infty} h_n^{\text{in}}(\mathbf{P}) \geq h_{\eta,\ell}$, and the claim follows by letting $\eta \rightarrow 1$ and $\ell \rightarrow \infty$.

It remains to show that $h_n^{\text{in}}(\mathbf{P}) \leq h^{\text{in}}(\mathbf{P})$. By (4.16) and the above, $h^{\text{in}}(\mathbf{P}) < \infty$. This implies that the kernel Φ admits a Radon-Nikodym density relative to Λ . Applying Remark 4.4 and Jensen's inequality as above, we conclude from (4.19) that

$$h_n^{\text{in}}(\mathbf{P}) \leq n^{-d} \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) \sum_{c \in T_s : c \subset [n]} \mathcal{H}(\Phi(s, \mathbf{T}_s, c, \cdot); \mathbb{1}_{\langle c \rangle} \Lambda).$$

The condition under the sum above implies that $m(c) \in [n]$. Using again (2.33) in its time-integrated version, we thus find that the last expression is not larger than $h^{\text{in}}(\mathbf{P})$. The proof is thus complete. \square

Remark 4.9. The inner entropy of a BRT \mathbf{P} in a window $W \in \mathbb{P}$ can be defined by considering the tessellations not only in W but also in some neighborhood of W . Namely, if Φ is a division kernel for \mathbf{P} and $r > 0$, one can introduce the quantity

$$\mathcal{H}_W^{r,\text{in}}(\mathbf{P}) = \int \mathbf{P}(d\mathbf{T}) \mathcal{H}(\mathbf{G}_W^{\Phi_{W+B_r}}(\cdot | \mathbf{T}_W^{W+B_r,0,\text{out}}); \mathbf{G}_W^\Lambda(\cdot | \mathbf{T}_W^{0,\partial})),$$

which is called the *inner entropy of \mathbf{P} in W with horizon r* . Here, the first of the conditional BRTs \mathbf{G} is as in Remark 4.5. A glance at the preceding proof then shows that Lemma 4.8 can be extended to yield

$$\lim_{n \rightarrow \infty} n^{-d} \mathcal{H}_{[n]}^{r,\text{in}}(\mathbf{P}) = h^{\text{in}}(\mathbf{P}).$$

Next we turn to the proof of Theorem 3.6, which is split into two lemmas.

Lemma 4.10. *The inner entropy density h^{in} is affine.*

Proof. As noticed after (4.15), relative entropy is a jointly convex function of probability measures. This shows that the inner entropies $\mathcal{H}_W^{\text{in}}(\mathbf{P}) = \mathcal{H}(\mathbf{P}_W; \mathbf{P}_W^{0,\partial} \otimes \mathbf{G}_W^\Lambda)$ are convex in \mathbf{P} , and so is their limit $h^{\text{in}}(\mathbf{P})$. The proof is therefore completed by showing that this limit is also concave. So, let $\mathbf{P}, \mathbf{P}' \in \mathcal{P}_\Theta$, $0 < a < 1$, $\hat{\mathbf{P}} = a\mathbf{P} + (1-a)\mathbf{P}'$, and assume without loss of generality that $h^{\text{in}}(\hat{\mathbf{P}}) < \infty$. By Lemma 4.8, it follows that $h_n^{\text{in}}(\hat{\mathbf{P}}) < \infty$ for all n . In particular, $\hat{\mathbf{P}}_{[n]} \ll \hat{\mathbf{P}}_{[n]}^{0,\partial} \otimes \mathbf{G}_{[n]}^\Lambda$ with a Radon-Nikodym density g_n . The Radon-Nikodym theorem further implies that $\mathbf{P}_{[n]} \ll \hat{\mathbf{P}}_{[n]}$ and $\mathbf{P}'_{[n]} \ll \hat{\mathbf{P}}_{[n]}$ with densities f_n and f'_n , respectively. It is clear that $af_n + (1-a)f'_n = 1$ almost surely for $\hat{\mathbf{P}}_{[n]}$. Moreover, it follows that $\mathbf{P}_{[n]}^{0,\partial} = f_n^{0,\partial} \hat{\mathbf{P}}_{[n]}^{0,\partial}$ for a suitable Radon-Nikodym density $f_n^{0,\partial}$. We conclude that $\mathbf{P}_{[n]} = (f_n g_n / f_n^{0,\partial}) \mathbf{P}_{[n]}^{0,\partial} \otimes \mathbf{G}_{[n]}^\Lambda$. Since $f_n \leq 1/a$ and $\int d\mathbf{P}_{[n]} \log f_n^{0,\partial} = \mathcal{H}(\mathbf{P}_{[n]}^{0,\partial}; \hat{\mathbf{P}}_{[n]}^{0,\partial}) \geq 0$, this gives

$$n^d h_n^{\text{in}}(\mathbf{P}) = \int d\mathbf{P}_{[n]} \log \frac{f_n g_n}{f_n^{0,\partial}} \leq \int d\mathbf{P}_{[n]} \log g_n + \log \frac{1}{a}.$$

Together with the analogous inequality for \mathbf{P}' , we finally end up with the estimate

$$a h_n^{\text{in}}(\mathbf{P}) + (1-a) h_n^{\text{in}}(\mathbf{P}') \leq n^{-d} \int d\hat{\mathbf{P}}_{[n]} \log g_n + o(1) = h_n^{\text{in}}(\hat{\mathbf{P}}) + o(1).$$

The result thus follows from Lemma 4.8 by letting $n \rightarrow \infty$. \square

As for the topological properties of h^{in} , we note first that its lower semicontinuity is a direct consequence of Lemma 4.8 and the lower semicontinuity of relative entropy; recall the discussion below (4.15). Since $\mathbf{T} \mapsto |T_{[1],1}|$ is the supremum of bounded local functions, it is also evident that the hitting intensity $i(\cdot)$ is lower semicontinuous. It follows that the restricted level sets $\mathcal{P}_{\Theta,P,\beta,\gamma}$ (as introduced in Theorem 3.6) are closed. The following lemma, which can be viewed as a refinement of Lemma 4.6, will imply that they are in fact compact; as the intensity bound is not needed here, we put $\beta = \infty$.

Lemma 4.11. *The restricted level sets $\mathcal{P}_{\Theta,P,\infty,\gamma}$ are locally equicontinuous in the following sense: For each $W \in \mathbb{P}$ and $0 \leq \gamma < \infty$ and every sequence $B_k \in \mathcal{B}_W$ with $B_k \downarrow \emptyset$ as $k \rightarrow \infty$, one has*

$$\lim_{k \rightarrow \infty} \sup_{\mathbf{P} \in \mathcal{P}_{\Theta,P,\infty,\gamma}} \mathbf{P}(B_k) = 0.$$

Proof. Let $W \in \mathbb{P}$ and a sequence $B_k \in \mathcal{B}_W$ with $B_k \downarrow \emptyset$ be given. Pick some $\varepsilon > 0$ and consider the events A_n defined in (4.17). By definition, $A_n \in \mathcal{B}_{[n],0}^\partial$. In particular, $\mathbf{P}(A_n)$ depends only on the initial distribution of \mathbf{P} , which is P for all $\mathbf{P} \in \mathcal{P}_{\Theta,P,\infty,\gamma}$. So there is an n with $\mathbf{P}(A_n) \geq 1 - \varepsilon$ for all $\mathbf{P} \in \mathcal{P}_{\Theta,P,\infty,\gamma}$.

Next, each $\mathbf{P} \in \mathcal{P}_{\Theta,P,\infty,\gamma}$ admits some division kernel Φ , and $\mathcal{H}_{[n]}^{\text{in}}(\mathbf{P}) \leq \eta := n^d \gamma$ by Lemma 4.8. Since

$$\mathcal{H}_{[n]}^{\text{in}}(\mathbf{P}) = \int \mathbf{P}(d\mathbf{T}) \mathcal{H}(\mathbf{G}_{[n]}^{\Phi_{[n]}}(\cdot | \mathbf{T}_{[n]}^{0,\partial}); \mathbf{G}_{[n]}^\Lambda(\cdot | \mathbf{T}_{[n]}^{0,\partial}))$$

by definition and Remark 4.5, we can conclude that the set

$$H_n := \{\mathbf{T} \in \mathbb{B}\mathbb{T} : \mathcal{H}(\mathbf{G}_{[n]}^{\Phi_{[n]}}(\cdot | \mathbf{T}_{[n]}^{0,\partial}); \mathbf{G}_{[n]}^\Lambda(\cdot | \mathbf{T}_{[n]}^{0,\partial})) \leq \eta/\varepsilon\}$$

in $\mathcal{B}_{[n]}^{0,\partial}$ has measure at least $1 - \varepsilon$ for \mathbf{P} . It follows that

$$\mathbf{P}(B_k) \leq 2\varepsilon + \mathbf{P}(B_k \cap A_n \cap H_n) = 2\varepsilon + \int_{A_n \cap H_n} \mathbf{P}(d\mathbf{T}) \mathbf{G}_{[n]}^{\Phi_{[n]}}(B_k | \mathbf{T}_{[n]}^{0,\partial})$$

because $\text{in}_{[n]}(s, \mathbf{T}_{[n]}^\partial) \supset W$ for all $\mathbf{T} \in A_n$ and all s ; recall (2.29). The next step is to use the inequality (4.15). For $\mathbf{T} \in A_n \cap H_n$, this inequality shows that

$$(\eta/\varepsilon^2) \mathbf{G}_{[n]}^{\Phi_{[n]}}(B_k | \mathbf{T}_{[n]}^\partial) \leq (\eta/\varepsilon) + \log \int d\Pi_W^\Lambda(T_{W,0}, \cdot) \exp[(\eta/\varepsilon^2) \mathbf{1}_{B_k}]$$

since $\mathbf{G}_{[n]}^\Lambda(\cdot | \mathbf{T}_{[n]}^\partial) = \Pi_W^\Lambda(T_{W,0}, \cdot)$ on \mathcal{B}_W when $\mathbf{T} \in A_n$. Inserting this into the previous inequality we find

$$\sup_{\mathbf{P} \in \mathcal{P}_{\Theta,P,\infty,\gamma}} \mathbf{P}(B_k) \leq 3\varepsilon + (\varepsilon^2/\eta) \int P_W(dT_W) \log \int d\Pi_W^\Lambda(T_W, \cdot) \exp[(\eta/\varepsilon^2) \mathbf{1}_{B_k}].$$

Letting $k \rightarrow \infty$, using the dominated convergence theorem, and noting that ε was chosen arbitrarily, we arrive at the lemma. \square

The preceding lemma verifies the conditions of Propositions (4.9) and (4.15) of [9], which imply that $\mathcal{P}_{\Theta,P,\infty,\gamma}$ is relatively compact and relatively sequentially compact within the class of all translation invariant BRTs. However, this does not yet imply that each limit of a net in $\mathcal{P}_{\Theta,P,\infty,\gamma}$ also satisfies the first-moment condition (2.31). This is the reason why we need to introduce the bound $i_1(\mathbf{P}) \leq \beta$ which makes this fact obvious. The proof of Theorem 3.6 is thus complete.

4.4 Free energy density, variational principle, existence

Throughout this section we fix a moderate division kernel Ψ . Our first item is the existence of the energy density.

Proof of Theorem 3.8. Let $\mathbf{P} \in \mathcal{P}_\Theta$ be a BRT with a covariant division kernel Φ . By Theorem 3.1(a), the inner energy of \mathbf{P} in a window $W \in \mathbb{P}$ can be written in the form

$$(4.21) \quad \mathcal{U}_W^{\text{in}}(\mathbf{P}; \Psi) = \int d\bar{\mathbf{P}}(s, \mathbf{T}_s) \sum_{c \in T_s : c \subset W} \int \Phi(s, \mathbf{T}_s, c, dH) \log \psi(s, \mathbf{T}_s, c, H).$$

Since Φ and ψ are covariant, the time-integrated version of the Palm formula (2.33) shows that the last term can be written in the form

$$\int d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \text{vol}(x : c + x \subset W) \int \Phi(s, \mathbf{T}_s, c, dH) \log \psi(s, \mathbf{T}_s, c, H).$$

Hence

$$n^{-d} \mathcal{U}_{[n]}^{\text{in}}(\mathbf{P}; \Psi) = u^{\text{in}}(\mathbf{P}; \Psi) + \delta_n(\mathbf{P}; \Psi)$$

with

$$|\delta_n(\mathbf{P}; \Psi)| \leq \kappa_\Psi \int d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \text{vol}(x \in [1] : c/n + x \notin [1]) \Phi(s, \mathbf{T}_s, c, \langle c \rangle)$$

by (M3). The volume term above is bounded by 1 and tends to 0 as $n \rightarrow \infty$. To apply the dominated convergence theorem we thus need to show that the total mass of $\bar{\mathbf{P}}^0 \otimes \Phi$ is finite. But the Palm formula and (4.8) show that this mass is at most $i_1(\mathbf{P})$. This completes the proof of the first part of Theorem 3.8 and implies the bound on $|u^{\text{in}}(\mathbf{P}; \Psi)|$.

The proof of the second part is similar: The Palm formula gives

$$\mathcal{V}_W^{\text{in}}(\mathbf{P}; \Psi) = \int d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \text{vol}(x : c + x \subset W) \int_{\langle c \rangle} \Lambda(dH) (\psi(s, \mathbf{T}_s, c, H) - 1)$$

and thus $n^{-d} \mathcal{V}_{[n]}^{\text{in}}(\mathbf{P}; \Psi) = v^{\text{in}}(\mathbf{P}; \Psi) + \delta'_n(\mathbf{P}; \Psi)$ with a remainder term δ'_n which, by assumption (M4), is bounded in modulus by κ'_Ψ times

$$\int d\bar{\mathbf{P}}^0(s, \mathbf{T}_s, c) \text{vol}(x \in [1] : c/n + x \notin [1]).$$

By (2.37) and the dominated convergence theorem, this bound vanishes in the limit $n \rightarrow \infty$. The proof of Theorem 3.8 is therefore complete. \square

Remark 4.12. Exploiting Theorem 3.1(a) and the Palm formula (2.33) in the same way as in the first part of the preceding proof, one finds that the energy density can be written in the alternative form

$$u^{\text{in}}(\mathbf{P}; \Psi) = \int \mathbf{P}(d\mathbf{T}) \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T}) : m(c) \in [1]} \log \psi(s, \mathbf{T}_{s-}, c, H),$$

in which the division kernel of \mathbf{P} does not appear. In particular, it follows that $u^{\text{in}}(\cdot; \Psi)$ is affine.

Turning to the proof of the variational principle, we introduce an *inner relative entropy of a BRT \mathbf{P} in a window $W \in \mathbb{P}$ with horizon $r = r_\Psi$ relative to Ψ* as follows: If \mathbf{P} admits a division kernel Φ we set, using the notation of Remark 4.5,

$$(4.22) \quad \mathcal{H}_W^{r,\text{in}}(\mathbf{P}; \Psi) = \int \mathbf{P}(d\mathbf{T}) \mathcal{H}(\mathbf{G}_W^{\Phi_{W+B_r}}(\cdot | \mathbf{T}_W^{W+B_r,0,\text{out}}); \mathbf{G}_W^{\Psi_{W+B_r}}(\cdot | \mathbf{T}_W^{W+B_r,0,\text{out}}));$$

otherwise we set $\mathcal{H}_W^{r,\text{in}}(\mathbf{P}; \Psi) = \infty$. (Compare this definition with Remark 4.9, where $\Psi = \Lambda^*$.) By the bounded-range property (M2) of Ψ , the conditional BRT $\mathbf{G}_W^{\Psi_{W+B_r}}$ in (4.22) actually coincides with \mathbf{G}_W^Ψ . We then have the following convergence to the quantity $h^{\text{in}}(\mathbf{P}; \Psi)$ in (3.11).

Corollary 4.13. *Let Ψ be a moderate division kernel and $r = r_\Psi$ its range. Then*

$$h^{\text{in}}(\mathbf{P}; \Psi) = \lim_{n \rightarrow \infty} n^{-d} \mathcal{H}_{[n]}^{r,\text{in}}(\mathbf{P}; \Psi)$$

for all $\mathbf{P} \in \mathcal{P}_\Theta$. The limit is finite if and only if $h^{\text{in}}(\mathbf{P}) < \infty$, and then Equation (3.12) holds.

Proof. An analog of Equation (4.20) gives for each n the identity

$$(4.23) \quad \mathcal{H}_{[n]}^{r,\text{in}}(\mathbf{P}; \Psi) = \mathcal{H}_{[n]}^{r,\text{in}}(\mathbf{P}) - \mathcal{U}_{[n]}^{\text{in}}(\mathbf{P}; \Psi) + \mathcal{V}_{[n]}^{\text{in}}(\mathbf{P}; \Psi),$$

which is a counterpart to (3.11). Also, the estimates in the proof of Theorem 3.8 show that the second and third term on the right-hand side are bounded in modulus by a finite constant times n^d . The convergence result thus follows directly from Remark 4.9 and Theorem 3.8 (together with Lemma 4.6 and Theorem 3.1).

Next, suppose that $h^{\text{in}}(\mathbf{P}) < \infty$ and let φ and ψ be the Radon-Nikodym densities of Φ and Ψ with respect to Λ^* . Inserting the explicit expressions for all quantities, we then obtain

$$\begin{aligned} h^{\text{in}}(\mathbf{P}) - u^{\text{in}}(\mathbf{P}; \Psi) + v^{\text{in}}(\mathbf{P}; \Psi) &= \int d\bar{\mathbf{P}}^0 \otimes \Lambda^* [-\varphi + \varphi \log \varphi - \varphi \log \psi + \psi] \\ &= \int d\bar{\mathbf{P}}^0 \otimes \Lambda^* \psi \varrho(\varphi/\psi) = \int d\bar{\mathbf{P}}^0 \mathcal{H}(\Phi; \Psi), \end{aligned}$$

which is (3.12). \square

The variational principle, Theorem 3.9, follows directly from Equation (3.12) and thus from the preceding corollary.

Next we address the existence problem for BRTs with given division kernel, as stated in Theorem 3.10. We still keep a moderate Ψ fixed and let $r = r_\Psi$ be its range. We also fix an initial distribution $P \in \mathcal{P}_\Theta(\mathbb{T})$. We will construct a BRT \mathbf{P} with initial distribution $\pi_0(\mathbf{P}) = P$ and division kernel Ψ as a cluster point of some approximating measures $\mathbf{P}^{n,\text{av}}$.

Specifically, for any n we let $\bar{n}+r$ and consider the shifted cubes $[n]_i = [n] + \bar{n}i$, $i \in \mathbb{Z}^d$, which are separated by a grid of corridors of width r . Let $[n]_\bullet = \bigcup_{i \in \mathbb{Z}^d} [n]_i$ be their union. We introduce a BRT \mathbf{P}^n for which the cells within the boxes $[n]_i$ evolve according to Ψ , while the cells that hit the corridors between these boxes evolve according to Λ^* . (This is inspired by the familiar construction of independent repetitions in disjoint blocks, but avoids an artificial cutting of cells at the block boundaries by using the STIT process in the corridors between the blocks.) Formally, we introduce the projection

$$\pi_{[n]_\bullet}^{0,\text{out}} : \mathbf{T} \mapsto \mathbf{T}_{[n]_\bullet}^{0,\text{out}} := \left(\bigcup_{i \in \mathbb{Z}^d} T_{[n]_i,0}^{\text{in}}, \bigcap_{i \in \mathbb{Z}^d} \mathbf{T}_{[n]_i}^{\text{out}} \right),$$

and define

$$(4.24) \quad \mathbf{P}^n = (P\Pi^\Lambda)_{[n]_\bullet}^{0,\text{out}} \otimes \bigotimes_{i \in \mathbb{Z}^d} \mathbf{G}_{[n]_i}^\Psi.$$

More explicitly, \mathbf{P}^n is defined by its integrals

$$\int f d\mathbf{P}^n = \int P\Pi^\Lambda(d\mathbf{T}) \prod_{i \in \mathbb{Z}^d} \int \mathbf{G}_{[n]_i}^\Psi(d\mathbf{s}_i | \mathbf{T}_{[n]_i}^{0,\text{out}}) f(\mathbf{T}_{[n]_\bullet}^{0,\text{out}} \cup \bigcup_i \mathbf{s}_i)$$

for measurable functions $f \geq 0$ on \mathbb{BT} . By the bounded-range property (M2), the conditional BRTs $\mathbf{G}_{[n]_i}^\Psi(\cdot | \mathbf{T}_{[n]_i}^{0,\text{out}})$ depend only on $\mathbf{T}_{[n]_\bullet}^{0,\text{out}}$, so that \mathbf{P}^n is well-defined. It is easily seen that \mathbf{P}^n is a BRT with initial distribution P and division kernel

$$(4.25) \quad \Psi^n(s, \mathbf{T}, c, \cdot) = \begin{cases} \Psi(s, \mathbf{T}, c, \cdot) & \text{if } c \subset [n]_i \text{ for some } i \in \mathbb{Z}^d, \\ \Lambda(\langle c \rangle \cap \cdot) & \text{otherwise.} \end{cases}$$

To achieve translation invariance, we introduce the average

$$(4.26) \quad \mathbf{P}^{n,\text{av}} = \bar{n}^{-d} \int_{[\bar{n}]} dx \mathbf{P}^n \circ \vartheta_x^{-1}.$$

The next two lemmas show that the BRTs $\mathbf{P}^{n,\text{av}}$ belong to a restricted level set of the inner entropy density, and thus have a cluster point.

Lemma 4.14. (a) *There exists a constant $\beta < \infty$ such that $i_1(\mathbf{P}^{n,\text{av}}) \leq \beta$ for all n .*

(b) *For every $\varepsilon > 0$ there exists some $\tau < \infty$ such that*

$$\int \mathbf{P}^{n,\text{av}}(d\mathbf{T}) |T_{[1],1}| \mathbf{1}_{\{|T_{[1],1}| \geq \tau\}} \leq \varepsilon \quad \text{for all } n.$$

Proof. Let $\kappa = \kappa_\Psi$. Since $\psi \leq e^\kappa$ by (M3), it follows that each kernel Ψ^n also has a Λ -density ψ^n satisfying $\psi^n \leq e^\kappa$ for all n . With the help of Remark 4.4, we can further conclude that this bound remains true after localisation to a window $W \in \mathbb{P}$ (relative to \mathbf{P}^n); that is, the localised kernel Ψ_W^n has a Λ -density ψ_W^n with $\psi_W^n \leq e^\kappa$. In particular, if $W = [1] + x$ is a translate of the unit cube, then

$$(4.27) \quad \Psi_W^n(s, \mathbf{T}_{W,s}, c, \langle c \rangle) \leq e^\kappa \Lambda(\langle [1] \rangle) =: \alpha < \infty$$

for all possible arguments. In view of Lemma 4.3, it follows that

$$(4.28) \quad \int \mathbf{P}^n(d\mathbf{T}) |T_{[1]+x,1}| \leq \beta := i_0(P) e^\alpha,$$

and statement (a) follows by averaging over x .

To prove (b) we still let $W = [1] + x$ and define $\varepsilon_1 = \varepsilon/4i_0(P)$. By Lemma 4.3, there exists a number τ_1 with

$$\sup_n \int \mathbf{P}^n(d\mathbf{T}) (|T_{W,1}| - \tau_1 |T_{W,0}|)_+ \leq \varepsilon_1 i_0(P) = \varepsilon/4.$$

For any τ_2 we then find (by distinguishing whether or not $\tau_1 |T_{W,0}| \leq \tau_2$) that

$$\int \mathbf{P}^n(d\mathbf{T}) (|T_{W,1}| - \tau_2)_+ \leq \varepsilon/4 + e^\alpha \int P(dT) |T_{[1]}| \mathbf{1}_{\{|T_{[1]}| > \tau_2/\tau_1\}},$$

which is at most $\varepsilon/2$ for suitable choice of τ_2 . Setting $\tau = 2\tau_2$ and using that

$$|T_{W,1}| \leq 2(|T_{W,1}| - \tau_2)_+ \quad \text{on} \quad \{|T_{W,1}| \geq \tau\},$$

we then see that

$$\int \mathbf{P}^n(d\mathbf{T}) |T_{[1]+x,1}| \mathbf{1}_{\{|T_{[1]+x,1}| \geq \tau\}} \leq \varepsilon$$

for all n . Statement (b) thus follows by taking the average over $x \in [\bar{n}]$. \square

Lemma 4.15. $h^{\text{in}}(\mathbf{P}^{n,\text{av}}; \Psi) \rightarrow 0$ as $n \rightarrow \infty$.

Proof. Fix any n and let \mathbf{P}^n and Ψ^n be given by (4.24) and (4.25). Consider the inner relative entropy of $\mathbf{P}^{n,\text{av}}$ relative to Ψ in a large cube $W = [k]$ and with horizon $r = r_\Psi$, as defined in (4.22). In concise notation, (4.22) reads

$$\mathcal{H}_W^{r,\text{in}}(\mathbf{P}; \Psi) = \mathcal{H}(\mathbf{P}_{W+B_r}; \mathbf{P}_W^{W+B_r,0,\text{out}} \otimes \mathbf{G}_W^\Psi).$$

As relative entropy is jointly measure convex, we have

$$\mathcal{H}_W^{r,\text{in}}(\mathbf{P}^{n,\text{av}}; \Psi) \leq \bar{n}^{-d} \int_{[\bar{n}]} dx \mathcal{H}_{W+x}^{r,\text{in}}(\mathbf{P}^n; \Psi).$$

To estimate this further we note that \mathbf{P}^n has the division kernel Ψ^n . A combination of (4.23), (4.21), (3.10) and an analog of (4.19) thus gives the formula

$$\mathcal{H}_{W+x}^{r,\text{in}}(\mathbf{P}^n; \Psi) = \int d\bar{\mathbf{P}}^n(s, \mathbf{T}_s) \sum_{c \in T_s: c \subset W+x} \mathcal{H}(\Psi^n; \Psi|_s, \mathbf{T}_s, c),$$

where $\mathcal{H}(\Psi^n; \Psi|_s, \mathbf{T}_s, c) = \mathcal{H}(\Psi^n(s, \mathbf{T}_s, c, \cdot); \Psi(s, \mathbf{T}_s, c, \cdot))$ for brevity. We can further use that $W+x \subset [k+\bar{n}]$ when $x \in [\bar{n}]$. Altogether, we obtain

$$\mathcal{H}_W^{r,\text{in}}(\mathbf{P}^{n,\text{av}}; \Psi) \leq \int d\bar{\mathbf{P}}^n(s, \mathbf{T}_s) \sum_{c \in T_s: c \subset [k+\bar{n}]} \mathcal{H}(\Psi^n; \Psi|_s, \mathbf{T}_s, c).$$

Next, it is clear from (4.25) that $\mathcal{H}(\Psi^n; \Psi| \cdot, \cdot, c) = 0$ when $c \subset [n]_i$ for some i . On the other hand, for any cell c hitting the corridors between the boxes $[n]_i$ we have $\mathcal{H}(\Psi^n; \Psi| \cdot, \cdot, c) = \mathcal{H}(\Lambda^*; \Psi| \cdot, \cdot, c)$, which is bounded by a constant. Indeed, the function $\varrho(a)$ defined in (3.5) is bounded by a multiple of $|a-1|$ as long as $a \leq e^{\kappa\Psi}$. Assumptions (M3) and (M4) therefore imply that

$$\mathcal{H}(\Lambda^*; \Psi| \cdot, \cdot, c) = \int_{\langle c \rangle} \Lambda(dH) \psi(\cdot, \cdot, c, H) \varrho(1/\psi(\cdot, \cdot, c, H)) \leq \tilde{\kappa}_\Psi$$

for some constant $\tilde{\kappa}_\Psi < \infty$ and all c hitting the corridors.

Now let $k = (\ell-1)\bar{n}$ for some integer ℓ and L_ℓ be such that $\{[1]+x : x \in L_\ell\}$ is a tessellation of $[\ell\bar{n}] \setminus [n]_\bullet$. The preceding estimates then show that

$$\begin{aligned} \mathcal{H}_{[k]}^{r,\text{in}}(\mathbf{P}^{n,\text{av}}; \Psi) &\leq \sum_{x \in L_\ell} \int d\bar{\mathbf{P}}^n(s, \mathbf{T}_s) \sum_{c \in T_s: c \cap ([1]+x) \neq \emptyset} \mathcal{H}(\Psi^n; \Psi|_s, \mathbf{T}_s, c) \\ &\leq \tilde{\kappa}_\Psi \sum_{x \in L_\ell} \int d\bar{\mathbf{P}}^n(s, \mathbf{T}_s) |T_{[1]+x,1}| \leq \tilde{\kappa}_\Psi \beta |L_\ell|. \end{aligned}$$

The last inequality comes from (4.28). Letting $\ell \rightarrow \infty$ and applying Corollary 4.13, we finally see that

$$h^{\text{in}}(\mathbf{P}^{n,\text{av}}; \Psi) \leq \tilde{\kappa}_\Psi \beta \lim_{\ell \rightarrow \infty} \frac{(\ell\bar{n})^d - \ell^d n^d}{((\ell-1)\bar{n})^d} = \tilde{\kappa}_\Psi \beta (1 - (n/\bar{n})^d).$$

This proves the lemma. \square

Combining Equation (3.11) with the last lemma and the bounds in Theorem 3.8 and Lemma 4.14(a), one finds that

$$h^{\text{in}}(\mathbf{P}^{n,\text{av}}) \leq (\kappa_\Psi + \kappa'_\Psi)\beta + 1 =: \gamma < \infty$$

when n is large enough. That is, the measures $\mathbf{P}^{n,\text{av}}$ eventually belong to the sequentially compact level set $\mathcal{P}_{\Theta, P, \beta, \gamma}$ of Theorem 3.6. This means that a subsequence converges in τ_{loc} to some \mathbf{P} in this set. We need to show that \mathbf{P} has the division kernel Ψ . In view of Theorem 3.9, this will follow once we have shown that $h^{\text{in}}(\mathbf{P}; \Psi) = 0$. By the last lemma, it is therefore sufficient to verify that $h^{\text{in}}(\cdot; \Psi)$ is lower semicontinuous on the closure of the sequence $\{\mathbf{P}^{n,\text{av}} : n \geq 1\}$. In view of Equation (3.11) and Theorem 3.6, this follows from the next lemma, which completes the proof of Theorem 3.10.

Lemma 4.16. *The functionals $u^{\text{in}}(\cdot; \Psi)$ and $v^{\text{in}}(\cdot; \Psi)$ are continuous on the closure \mathcal{C} of the sequence $\{\mathbf{P}^{n,\text{av}} : n \geq 1\}$.*

Proof. First we observe that the estimate in Lemma 4.14(b) holds not only for all $\mathbf{P}^{n,\text{av}}$, but even for all $\mathbf{P} \in \mathcal{C}$. This is because the integral there is a lower semicontinuous function of the integrating measure. We further know from Lemma 4.6 that each $\mathbf{P} \in \mathcal{C}$ satisfies (LAC). Hence, Theorem 3.8 and Remark 4.12 can be applied.

It follows that $u^{\text{in}}(\mathbf{P}; \Psi) = \int u \, d\mathbf{P}$ for the function

$$u(\mathbf{T}) = \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T}): m(c) \in [1]} \log \psi(s, \mathbf{T}_{s-}, c, H)$$

on \mathbb{BT} , which in general is neither bounded nor local. We will therefore replace u by a truncated version

$$u_{\tau,\ell}(\mathbf{T}) = \mathbb{1}_{\{|\mathbf{T}_{[1],1}| \leq \tau\}} \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T}): m(c) \in [1], r(c) \leq \ell} \log \psi(s, \mathbf{T}_{s-}, c, H),$$

for suitable numbers τ and ℓ ; $r(c)$ is again the radius of c . The function $u_{\tau,\ell}$ is bounded in modulus by $\kappa_{\Psi}\tau$ and also local because of (M2). It differs from u by at most $\kappa_{\Psi}(\delta_{\tau} + \delta_{\ell})$ with the error functions

$$\delta_{\tau}(\mathbf{T}) = \mathbb{1}_{\{|\mathbf{T}_{[1],1}| > \tau\}} |\mathbf{T}_{[1],1}|, \quad \delta_{\ell}(\mathbf{T}) = \sum_{(s,c,H) \in \mathcal{D}(\mathbf{T}): m(c) \in [1]} \mathbb{1}_{\{r(c) > \ell\}}.$$

As noticed at the beginning of this proof, we have $\sup_{\mathbf{P} \in \mathcal{C}} \int \delta_{\tau} \, d\mathbf{P} \rightarrow 0$ as $\tau \rightarrow \infty$. On the other hand, the function δ_{ℓ} is not larger than

$$\delta'_{\ell}(\mathbf{T}) = \sum_{c_0 \in T_0: c_0 \cap [1] \neq \emptyset, r(c_0) > \ell} |T_{c_0 \cap [1],1}|,$$

and Lemma 4.3 gives the estimate

$$\sup_{\mathbf{P} \in \mathcal{C}} \int \delta'_{\ell} \, d\mathbf{P} \leq e^{\alpha} \int P(dT) \sum_{c_0 \in T_0: c_0 \cap [1] \neq \emptyset} \mathbb{1}_{\{r(c_0) > \ell\}}$$

for the constant α in (4.27) because each $\mathbf{P} \in \mathcal{C}$ has initial distribution P . This bound does not depend on n and tends to 0 as $\ell \rightarrow \infty$ because $i_0(P) < \infty$. We have thus shown that the restriction of $u^{\text{in}}(\cdot; \Psi)$ to \mathcal{C} is the uniform limit of the functions $\mathbf{P} \mapsto \int u_{\tau,\ell} \, d\mathbf{P}$, which are continuous in τ_{loc} .

The analogous result for $v^{\text{in}}(\mathbf{P}^{n,\text{av}}; \Psi)$ is achieved in a similar way by truncating the function

$$v(\mathbf{T}) = \int_0^1 ds \sum_{c \in T_s: m(c) \in [1]} \int_{\langle c \rangle} \Lambda(dH) (\psi(s, \mathbf{T}_s, c, H) - 1)$$

and using (M4). □

As the proof of Theorem 3.10 is now complete, we turn to its corollary.

Proof of Corollary 3.11. Suppose $\mathbf{P} \in \mathcal{G}_{\Theta}(\Psi)$ is not extremal in \mathcal{P}_{Θ} . Then $\mathbf{P} = a\mathbf{P}^1 + (1-a)\mathbf{P}^2$ for some $0 < a < 1$ and two distinct BRTs $\mathbf{P}^1, \mathbf{P}^2 \in \mathcal{P}_{\Theta}$. By Theorem 3.6, Remark 4.12 and Theorem 3.9, it follows that

$$0 = h^{\text{in}}(\mathbf{P}; \Psi) = a h^{\text{in}}(\mathbf{P}^1; \Psi) + (1-a) h^{\text{in}}(\mathbf{P}^2; \Psi),$$

so that $\mathbf{P}^1, \mathbf{P}^2$ both belong to $\mathcal{G}_{\Theta}(\Psi)$. Hence, \mathbf{P} is not extremal in $\mathcal{G}_{\Theta}(\Psi)$. □

Our final observation concerns the uniqueness problem discussed in Remark 3.12. We will exploit the fact that, in one space dimension, we always have that $\sum_{c \in T_W} \Lambda(\langle c \rangle) = \Lambda(\langle W \rangle)$ when $W \in \mathbb{P}$ and $T_W \in \mathbb{T}_W$. Consider the following variants of conditions (M3) and (M4):

(M3') Ψ is *STIT-bounded*, in that $\Psi \leq K_\Psi \Lambda^*$ for some constant $K_\Psi < \infty$.

(M4') Ψ is *STIT for large cells*, in that $\Psi(\cdot, \cdot, c, \cdot) = \Lambda(\langle c \rangle \cap \cdot)$ whenever $\text{diam}(c) \geq r'_\Psi$ for some constant $r'_\Psi < \infty$.

Proposition 4.17. *Suppose that the space dimension is $d = 1$. Let $P \in \mathcal{P}(\mathbb{T})$ and Ψ be a division kernel satisfying (M2), (M3') and (M4'). Then there exists at most one BRT for Ψ with initial distribution P .*

Proof. Suppose there exist two distinct BRTs \mathbf{P}, \mathbf{P}' for Ψ with the same initial distribution P . Consider the difference measure $\mathbf{P}^\delta = \mathbf{P} - \mathbf{P}'$ and fix an interval $[k] \in \mathbb{P}$. Let g be $\mathcal{B}_{[k]}$ -measurable with $|g| \leq 1$. Using property (c) of Theorem 3.1, we obtain for each $0 < t \leq 1$ the identity

$$\int g \, d\mathbf{P}_t^\delta = \int_0^t ds \int \mathbf{P}_s^\delta(d\mathbf{T}_s) \mathbb{L}_s^\Psi g(\mathbf{T}_s)$$

with

$$\mathbb{L}_s^\Psi g(\mathbf{T}_s) = \sum_{c \in T_s: c \cap [k] \neq \emptyset} \int_{\langle c \cap [k] \rangle} \Psi(s, \mathbf{T}_s, c, dH) [g(\mathcal{O}_{s,c,H}(\mathbf{T}_s)) - g(\mathbf{T}_s)].$$

Now, (M2) and (M4') imply that $\mathbb{L}_s^\Psi g(\mathbf{T}_s)$ depends only on $\mathbf{T}_{[k+r],s}$ with $r = 2(r_\Psi + r'_\Psi)$. On the other hand, using (M3') and the additivity of $c \mapsto \Lambda(\langle c \rangle)$ we find that

$$|\mathbb{L}_s^\Psi g(\mathbf{T}_s)| \leq 2K_\Psi \sum_{c \ni T_s: c \cap [k] \neq \emptyset} \Lambda(\langle c \cap [k] \rangle) = 2K_\Psi \Lambda(\langle [k] \rangle) =: \alpha k.$$

The total variation norm $\delta_k(t) := \|\mathbf{P}_{[k],t}^\delta\|$ thus satisfies the inequality

$$(4.29) \quad \delta_k(t) \leq \alpha k \int_0^t \delta_{k+r}(s) \, ds$$

of Gronwall type. (Note that δ_k is increasing and therefore measurable.) Since $\delta_{k+nr}(s) \leq 2$, we obtain by n -fold iteration

$$\delta_k(t) \leq 2\alpha^n (k + nr)^n t^n / n! \leq 2e^k (\alpha t e^r)^n$$

and thus, in the limit as $n \rightarrow \infty$, $\delta_k(t) = 0$ for all $t < \varepsilon := 1/(\alpha e^r)$ and all k . Inserting this into (4.29) and repeating the estimate, we obtain that $\delta_k(t) = 0$ for all $t < 2\varepsilon$ and all k . Continuing in this way we finally find that $\delta_k(1) = 0$ for all k , which means that $\mathbf{P} = \mathbf{P}'$. \square

References

- [1] ARAK, T.; SURGAILIS, D.: *Markov fields with polygonal realizations*, Probab. Theory Related Fields **80**, 543–579 (1989).
- [2] ARAK, T.; SURGAILIS, D.: *Consistent polygonal fields*, Probab. Theory Related Fields **89**, 319–346 (1991).
- [3] BERTIN, E.; BILLIOT, J.-M.; DROUILHET, R.: *Existence of Delaunay pairwise Gibbs point process with superstable component*, J. Stat. Phys. **95**, 719–744 (1999).
- [4] BERTOIN, J.: *Random Fragmentation and Coagulation Processes*, Cambridge University Press (2006).

- [5] DEREUDRE, D.; DROUILHET, R.; GEORGII, H.-O.: *Existence of Gibbsian point processes with geometry-dependent interactions*, Probab. Theory Related Fields **153**, 643–670 (2012).
- [6] DEREUDRE, D.; GEORGII, H.-O.: *Variational characterisation of Gibbs measures with Delaunay triangle interaction*, Electron. J. Probab. **14**, 2438–2462 (2009).
- [7] FELLER, W.: *The integrodifferential equations of purely discontinuous Markoff processes*, Trans. Amer. Math. Soc. **48**, 488–515 (1940).
- [8] FÖLLMER, H.; SNELL, J.L.: *An “inner” variational principle for Markov fields on a graph*, Z. Wahrscheinlichkeitstheorie verw. Gebiete **39**, 187–195 (1977).
- [9] GEORGII, H.-O.: *Gibbs Measures and Phase Transitions*, 2nd Edition, de Gruyter, Berlin (2011).
- [10] KALLENBERG, O.: *Random Measures*, 3rd Edition, Akademie-Verlag, Berlin (1983).
- [11] KALLENBERG, O.: *Foundations of Modern Probability*, 2nd Edition, Springer, New York (2002).
- [12] MECKE, J.; NAGEL, W.; WEISS, V.: *A global construction of homogeneous random planar tessellations that are stable under iteration*, Stochastics **80**, 51–67 (2008).
- [13] NAGEL, W.; WEISS, V.: *Crack STIT tessellations: characterization of stationary random tessellations stable with respect to iteration*, Adv. Appl. Probab. **37**, 859–883 (2005).
- [14] OKABE, A.; BOOTS, B.; SUGIHARA, K.; CHIU, S.N.: *Spatial Tessellations — Concepts and Applications of Voronoi Diagrams*, Wiley, Chichester (2000).
- [15] REDENBACH, C.; THÄLE, C.: *On the arrangement of cells in planar STIT and Poisson line tessellations*, Methodol. Comput. Appl. Probab. DOI: 10.1007/s11009-011-9273-4 (2011+).
- [16] ROSS, S.M.: *Introduction to Probability Models*, 8th ed., Academic Press, Amsterdam (2003).
- [17] SCHNEIDER, R.; WEIL, W.: *Stochastic and Integral Geometry*, Springer, Berlin (2008).
- [18] SCHREIBER, T.; THÄLE, C.: *Second-order properties and central limit theory for the vertex process of iteration infinitely divisible and iteration stable random tessellations in the plane*, Adv. Appl. Probab. **42**, 913–935 (2010).
- [19] SCHREIBER, T.; THÄLE, C.: *Intrinsic Volumes of the maximal polytope process in higher dimensional STIT tessellations*, Stoch. Proc. Appl. **121**, 989–1012 (2011).
- [20] SCHREIBER, T.; THÄLE, C.: *Shape-driven Markov tessellations*, Stochastics DOI: 10.1080/17442508.2011.654344 (2012+).
- [21] SCHREIBER, T.; THÄLE, C.: *Limit theorems for iteration stable tessellations*, accepted for publication in The Annals of Probability (2011+).
- [22] SCHREIBER, T.; THÄLE, C.: *Geometry of iteration stable tessellations: Connection with Poisson hyperplanes*, accepted for publication in Bernoulli (2012+).
- [23] SCHREIBER, T.; THÄLE, C.: *Second-order theory for iteration stable tessellations*, Probab. Math. Stat. **32**, 281–300 (2012).
- [24] STOYAN, D.; KENDALL, D.G.; MECKE, J.: *Stochastic Geometry*, 2nd Edition, Wiley, Chichester (1995).
- [25] THÄLE, C.; WEISS, V.; NAGEL, W.: *Spatial STIT tessellations: distributional results for I -segments*, Adv. Appl. Probab. **44**, 635–654 (2012).
- [26] VARADHAN, S.R.S.: *Large deviations and applications*, Lect. Notes Math. **1362**, 1–49, Springer, Berlin (1988).